



# Using Aggregate Relocation to Upgrade Controller Hardware Running ONTAP 9.5 or Later

November 2018 | 215-13625\_AO\_UR001  
doccomments@netapp.com



# Contents

<b>Using aggregate relocation to upgrade controller hardware .....</b>	<b>5</b>
Automating the controller upgrade process .....	5
<b>Deciding whether to use the aggregate relocation guide .....</b>	<b>7</b>
<b>Required tools and documentation .....</b>	<b>8</b>
Guidelines for upgrading controllers with ARL .....	9
<b>Overview of the ARL upgrade .....</b>	<b>10</b>
<b>Upgrading the node pair .....</b>	<b>12</b>
<b>Stage 1 Preparing for upgrade .....</b>	<b>13</b>
Preparing the nodes for upgrade .....	13
Correcting aggregate ownership if an ARL precheck fails .....	16
Licensing .....	16
Getting an IP address of an external key management server for storage encryption .....	17
Managing authentication using KMIP servers .....	17
Managing authentication using an onboard key manager .....	18
Quiescing the SnapMirror relationships (optional) .....	18
<b>Stage 2 Relocating and retiring node1 .....</b>	<b>19</b>
Relocating non-root aggregates and NAS data LIFs owned by node1 to node2 .....	19
Relocating failed or vetoed aggregates .....	21
Retiring node1 .....	21
Preparing for netboot .....	22
<b>Stage 3 Installing and booting node3 .....</b>	<b>23</b>
Installing and booting node3 .....	23
Setting the FC or UTA/UTA2 configuration on node3 .....	27
Configuring FC ports on node3 .....	27
Checking and configuring UTA/UTA2 ports on node3 .....	29
Mapping ports from node1 to node3 .....	34
Joining the quorum when a node has a different set of network ports .....	37
Verifying the node3 installation .....	39
Re-creating VLANs, ifgrps, and broadcast domains on node3 .....	39
Moving non-root aggregates and NAS data LIFs owned by node1 from node2 to node3 .....	40
<b>Stage 4 Relocating and retiring node2 .....</b>	<b>43</b>
Relocating non-root aggregates and NAS data LIFs from node2 to node3 .....	43
Retiring node2 .....	45
<b>Stage 5 Installing and booting node4 .....</b>	<b>46</b>
Installing and booting node4 .....	46
Setting the FC or UTA/UTA2 configuration on node4 .....	50
Configuring FC ports on node4 .....	50
Checking and configuring UTA/UTA2 ports on node4 .....	52
Mapping ports from node2 to node4 .....	56

Joining the quorum when a node has a different set of network ports .....	60
Verifying the node4 installation .....	61
Re-creating VLANs, ifgrps, and broadcast domains on node4 .....	62
Moving non-root aggregates and NAS data LIFs owned by node2 from node3 to node4 .....	62
<b>Stage 6 Completing the upgrade .....</b>	<b>66</b>
Ensuring that the new controllers are set up correctly .....	66
Setting up Storage Encryption on the new controller module .....	68
Decommissioning the old system .....	69
Resuming SnapMirror operations .....	69
<b>Troubleshooting .....</b>	<b>71</b>
Aggregate relocation failures .....	71
Checking for aggregate relocation failure .....	71
Aggregates originally on node1 are owned by node4 after completion of the upgrade .....	71
Reboots, panics, or power cycles .....	72
Reboots, panics, or power cycles during the pre-check phase .....	72
Reboots, panics, or power cycles during first resource-release phase .....	72
Reboots, panics, or power cycles during the first verification phase .....	74
Reboots, panics, or power cycles during first resource-regain phase .....	74
Reboots, panics, or power cycles during post-check phase .....	75
Reboots, panics, or power cycles during second resource-release phase .....	75
Reboots, panics, or power cycles during the second verification phase .....	75
Issues that can arise in multiple stages of the procedure .....	76
Unexpected storage failover show command output .....	76
LIF migration failure .....	77
<b>Copyright information .....</b>	<b>78</b>
<b>Trademark information .....</b>	<b>79</b>
<b>How to send comments about documentation and receive update notifications .....</b>	<b>80</b>
<b>Index .....</b>	<b>81</b>

## Using aggregate relocation to upgrade controller hardware

---

You can nondisruptively upgrade the controller hardware on a pair of nodes running ONTAP 9.5 by migrating non-root aggregates from the original nodes to the new nodes in the same cluster. The data hosted on the nodes that are being upgraded is accessible during the upgrade.

### About this task

This procedure uses a method called aggregate relocation (ARL), which takes advantage of the HA configuration to enable you to move ownership of non-root aggregates from one node to another, if they share storage within the same cluster.

During the procedure, you upgrade the original controller hardware with the replacement controller hardware, relocating the ownership of non-root aggregates. You migrate aggregates multiple times from node to node to ensure that at least one node is serving data from the aggregates throughout the upgrade procedure. You also migrate data logical interfaces (LIFs) and assign the network ports on the new controller to the interface groups as you proceed.

**Important:** This procedure is complex and assumes that you have advanced ONTAP administration skills. You also should read and understand the [Guidelines for upgrading controllers with ARL](#) on page 9 and [Overview of the ARL upgrade](#) on page 10 sections before beginning the upgrade.

In this document, the original nodes are called *node1* and *node2*, and the new nodes are called *node3* and *node4*. During the described procedure, *node1* is replaced by *node3*, and *node2* is replaced by *node4*.

**Attention:** This procedure assumes that the replacement controller hardware is new and has not been used. The steps required to prepare used controllers with the `wipeconfig` command are not included in this procedure. You must contact technical support if the replacement controller hardware was previously used, especially if the controllers were running Data ONTAP in 7-Mode.

The terms *node1*, *node2*, *node3*, and *node4* are used only to distinguish between the original and new nodes. When following the procedure, you must substitute the real names of your original and new nodes. However, in reality, the names of the nodes do not change: *node3* has the name *node1*, and *node4* has the name *node2* after the controller hardware is upgraded.

**Note:** You can use this procedure to upgrade the controller hardware in clusters with more than two nodes; however, you need to perform the procedure separately for each HA pair in the cluster.

This procedure applies to FAS systems, V-Series systems, AFF systems, and systems with FlexArray Virtualization Software.

FAS systems released after ONTAP 9.5 can attach to storage arrays if the required license is installed. The existing V-Series systems are supported in ONTAP 9.5. See the V-Series Support Matrix at [Hardware Universe](#) at [hww.netapp.com](http://hww.netapp.com) for information about the storage array and V-Series models.

**Note:** This document uses the term *systems with FlexArray Virtualization Software* to refer to systems that belong to these new platforms. It uses the term *V-Series system* to refer to the separate hardware systems that can attach to storage arrays.

## Automating the controller upgrade process

During a controller update, the controller is replaced with another controller running a newer or more powerful platform. Earlier versions of this guide contained instructions for a nondisruptive controller

## 6 | Using Aggregate Relocation to Upgrade Controller Hardware Running ONTAP 9.5 or Later

update process that was comprised of entirely manual steps. This guide provides the steps for the new automated procedure.

The manual process was lengthy and complex but in this simplified procedure you can implement a controller update using aggregate relocation, allowing for more efficient nondisruptive upgrades for HA pairs. There are significantly fewer manual steps, especially around validation, collection of information, and post checks.

## Deciding whether to use the aggregate relocation guide

---

This guide describes how to upgrade the storage controllers in an HA pair with new controllers while keeping all the existing data and disks. This is a complex procedure that should be used only by experienced administrators.

You should use this guide under the following circumstances:

- You do not want to add the new controllers as a new HA pair to the cluster and migrate the data using volume moves.
- You are experienced in administering ONTAP and are comfortable with the risks of working in the diagnostic privilege mode.
- You do not have a MetroCluster configuration.

The aggregate relocation procedure for ONTAP 9.5 does not support MetroCluster.

If you have a system that uses MetroCluster, AFF80x0, AFF800, see *Using Aggregate Relocation to Upgrade Controller Hardware Running ONTAP 9.0 to 9.4* for detailed information about upgrading ONTAP and relocating aggregates on those systems.

**Note:** You can use NetApp Storage Encryption (NSE) with this procedure. NetApp Volume Encryption (NVE) does not apply.

If you prefer a different method of upgrading the controller hardware and are willing to do volume moves, use the *Controller Hardware Upgrade Express Guide* instead. See the ONTAP 9 Documentation Center to access ONTAP 9 product documentation.

### Related information

[Controller Hardware Upgrade Express Guide](#)  
[ONTAP 9 Documentation Center](#)

## Required tools and documentation

---

You must have specific tools to install the new hardware, and you need to reference other documents during the upgrade process. You also need to record information essential to completing the controller upgrade; a worksheet is provided to record information.

You need the following tools to perform the upgrade:

- Grounding strap
- #2 Phillips screwdriver

Document	Contents
<i>ONTAP 9 Software Setup Guide</i>	Describes how to set up and configure NetApp systems
<i>ONTAP 9 System Administration Reference</i>	Describes how to administer ONTAP systems, shows you how to use the CLI interface, how to access the cluster, how to manage nodes, and much more
<i>ONTAP 9 Upgrade and Revert/Downgrade Guide</i>	Contains instructions for downloading and upgrading ONTAP
<i>ONTAP 9 Disks and Aggregates Power Guide</i>	Describes how to manage ONTAP physical storage, shows you how to create, expand, and manage aggregates, how to work with Flash Pool aggregates, how to manage disks, and how to manage RAID policies
<i>ONTAP 9 Logical Storage Management Guide</i>	Describes how to efficiently manage your logical storage resources, using volumes, FlexClone volumes, files and LUNs, FlexCache volumes, deduplication, compression, qtrees, and quotas
<i>Installation and Setup Instructions for the model of the new nodes</i>	Contains instructions for installing and cabling the new system.
<i>FlexArray Virtualization Installation Requirements and Reference Guide</i>	Contains cabling instructions and other information for FlexArray Virtualization systems
<i>Hardware Universe at <a href="http://hwu.netapp.com">hwu.netapp.com</a></i>	Contains information about system LEDs and system messages to aid troubleshooting
The appropriate disk shelf guide	Contains instructions for installing and monitoring disk shelves and replacing disk shelf devices
<i>ONTAP 9 SAN Administration Guide</i>	Describes how to configure and manage iSCSI and FC protocols for SAN environments
<i>ONTAP 9 SAN Configuration Guide</i>	Contains information about FC and iSCSI topologies and wiring schemes

The NetApp Support Site also contains documentation about NICs and other hardware that you might use with your system. It also contains the *Hardware Universe*, which provides information about the hardware that the new system supports. See the ONTAP 9 Documentation Center to access ONTAP 9 product documentation.

### Related information

[ONTAP 9 Documentation Center](#)

## Guidelines for upgrading controllers with ARL

To understand whether you can use aggregate relocation (ARL) to upgrade a pair of controllers running ONTAP 9.5 and later depends on the platform and the configuration of both the original and replacement controllers.

### Supported upgrades for ARL

When you upgrade a pair of nodes using this ARL procedure for ONTAP 9.5 and later, you must verify that ARL can be performed on the original and replacement controllers.

You should check the size of all defined aggregates and number of disks supported by the original system. You must then compare the aggregate sizes and number of disks supported to the aggregate size and number of disks supported by the new system. This information is available in the *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com). The aggregate size and the number of disks supported by the new system must be equal to or greater than the aggregate size and number of disks supported by the original system.

You should validate in the cluster mixing rules whether new nodes can become part of the cluster with the existing nodes, when the original controller is replaced. For more information about cluster mixing rules, see the *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com).

### Two-node switchless clusters

If you are upgrading nodes in a two-node switchless cluster, you can leave the nodes in the switchless cluster while performing the upgrade. You do not need to convert them to a switched cluster.

### Troubleshooting

If any problems occur while upgrading the controllers, you can refer to the [Troubleshooting](#) on page 71 section at the end of the procedure for more information and possible solutions.

If you do not find a solution to the problem you encountered, you should contact technical support.

### Related information

[Controller Hardware Upgrade Express Guide](#)

## Overview of the ARL upgrade

---

Before you upgrade the nodes using ARL, you should understand how the procedure works. In this document, the procedure is broken down into several stages.

During the procedure, you upgrade the original controller hardware with the replacement controller hardware, one controller at a time, taking advantage of the HA configuration to relocate the ownership of non-root aggregates. All non-root aggregates must undergo two relocations to reach their final destination, which is the correct upgraded node.

Each aggregate has a home owner and current owner. The home owner is the actual owner of the aggregate, and the current owner is the temporary owner.

<b>1</b>	<p>Stage 1: Preparing for the upgrade</p> <ol style="list-style-type: none"> <li>1. Run prechecks to verify that the operation can be performed.</li> <li>2. Prepare for netbooting (optional).</li> </ol> <p>Aggregate ownership at the end of Stage 1:</p> <ul style="list-style-type: none"> <li>• Node1 is the home owner and current owner of the node1 aggregates.</li> <li>• Node2 is the home owner and current owner of the node2 aggregates.</li> </ul>
<b>2</b>	<p>Stage 2: Relocating and retiring node1</p> <ol style="list-style-type: none"> <li>1. Move non-root aggregates from node1 to node2.</li> <li>2. Move non-SAN data LIFs owned by node1 to node2.</li> <li>3. Retire node1.</li> </ol> <p>Aggregate ownership at the end of Stage 2:</p> <ul style="list-style-type: none"> <li>• Node2 is the current owner of node1 aggregates.</li> <li>• Node2 is the home owner and current owner of node2 aggregates.</li> </ul>
<b>3</b>	<p>Stage 3: Installing and booting node3</p> <ol style="list-style-type: none"> <li>1. Install and boot node3.</li> <li>2. Set the FC or CNA configuration.</li> <li>3. Map ports from node1 to node3.</li> <li>4. Verify node3 is successfully installed.</li> <li>5. Move non-SAN data LIFs owned by node1 from node2 to node3.</li> <li>6. Relocate non-root aggregates owned by node1 from node2 to node3.</li> </ol> <p>Aggregate ownership at the end of Stage 3:</p> <ul style="list-style-type: none"> <li>• Node3 is the home owner and current owner of node1 aggregates.</li> <li>• Node2 is the home owner and current owner of node2 aggregates.</li> </ul>

<p style="text-align: center;"><b>4</b></p>	<p>Stage 4: Relocating and retiring node2</p> <ol style="list-style-type: none"> <li>1. Relocate non-root aggregates from node2 to node3.</li> <li>2. Move non-SAN data LIFs owned by node2 to node 3.</li> <li>3. Retire node2.</li> </ol> <p>Aggregate ownership at the end of Stage 4:</p> <ul style="list-style-type: none"> <li>• Node3 is the home owner and current owner of aggregates that originally belonged to node1.</li> <li>• Node2 is the home owner of node2 aggregates.</li> <li>• Node3 is the current owner of node2 aggregates.</li> </ul>
<p style="text-align: center;"><b>5</b></p>	<p>Stage 5: Installing and booting node4</p> <ol style="list-style-type: none"> <li>1. Install and boot node4.</li> <li>2. Set the FC or CNA configuration on node4.</li> <li>3. Map ports from node2 to node4.</li> <li>4. Verify the node4 installation.</li> <li>5. Relocate non-SAN data LIFs owned by node2 from node3 to node4 and verify the SAN LIFs on node4.</li> </ol> <p>Aggregate ownership at the end of Stage 5:</p> <ul style="list-style-type: none"> <li>• Node3 is the home owner and current owner of the aggregates that originally belonged to node1.</li> <li>• Node4 is the home owner and current owner of aggregates that originally belonged to node2.</li> </ul>
<p style="text-align: center;"><b>6</b></p>	<p>Stage 6: Completing the upgrade</p> <ol style="list-style-type: none"> <li>1. Verify the system is set up correctly.</li> <li>2. Set up Storage Encryption on the new controller module.</li> <li>3. Resume SnapMirror operations if needed.</li> </ol>

## **Upgrading the node pair**

---

To upgrade the node pair, you need to prepare the original nodes and then perform a series of steps on both the original and new nodes. You can then decommission the original nodes.

## Stage 1 Preparing for upgrade

---

During Stage 1, you must prepare the nodes for the upgrade and run a series of prechecks. You might need to rekey disks for Storage Encryption. You must also prepare to netboot the new controllers.

### Steps

1. [Preparing the nodes for upgrade](#) on page 13
2. [Getting an IP address of an external key management server for storage encryption](#) on page 17
3. [Quiescing the SnapMirror relationships \(optional\)](#) on page 18

## Preparing the nodes for upgrade

### Steps

1. Begin the controller replacement process by entering the following command on the ONTAP command line:

```
system controller replace start -nodes node_names
```

**Note:** This command can only be executed at the advanced privilege level:

```
set -privilege advanced
```

You will see the following output:

```
Warning: 1. Current Data ONTAP version is 9.5
          Before starting controller replacement operation, ensure that the new
          controllers are running the version 9.5

          2. Verify that NVMEM or NVRAM batteries of the new controllers are
          charged, and charge them if they are not. You need to physically
          check the new controllers to see if the NVMEM or NVRAM batteries are
          charged. You can check the battery status either by connecting to a
          console or using SSH, logging into the Service Processor (SP) for
          your system, and use the system sensors to see if the battery has a
          sufficient charge.

          Attention: Do not try to clear the NVRAM contents. If there is a need
          to clear the contents of NVRAM, contact NetApp technical support.

          3. If you are replacing the controllers with a used one, ensure that
          you run a wipeconfig before controller replacement operation.

Do you want to continue? {y|n}: y
```

2. Press y.

You will see the following output:

```
Controller replacement operation: Prechecks in progress..
Controller replacement operation has been paused for user
intervention.
```

The system runs the following prechecks; record the output of each precheck for use later in the procedure:

Precheck	Description
Cluster Health Check	Checks all the nodes in the cluster to confirm they are healthy.

Precheck	Description
MCC Cluster Check	Checks if the system is a MetroCluster configuration. MetroCluster configurations are not supported by this procedure.
Aggregate Relocation Status Check	Checks whether an aggregate relocation is already in progress. If another aggregate relocation is in progress, the check fails.
Model Name Check	Checks whether the controller models are supported for this procedure. If the models are not supported, the task fails.
Cluster Quorum Check	Checks that the nodes being replaced are in quorum. If the nodes are not in quorum, the task fails.
Image Version Check	Checks that the nodes being replaced run the same version of ONTAP. If the ONTAP image versions are different, the task fails.  The new nodes must have the same version of ONTAP 9.x installed on them that is installed on the original nodes. If the new nodes have a different version of ONTAP installed, you need to netboot the new controllers after you install them. See the <i>ONTAP 9 Upgrade and Revert/Downgrade Guide</i> for instructions to upgrade ONTAP.
HA Status Check	Checks if both the nodes being replaced are in a high-availability (HA) pair configuration. If storage failover is not enabled for the controllers, the task fails.
Aggregate Status Check	If the nodes being replaced own aggregates for which they are not the home owner, the task fails.  The nodes should not own any non-local aggregates.
Disk Status Check	If any nodes being replaced have missing or failed disks, the task fails.  If any disks are missing, see the <i>ONTAP 9 Disks and Aggregates Power Guide</i> , the <i>ONTAP 9 Logical Storage Management Guide</i> , and the <i>ONTAP 9 High-Availability Configuration Guide</i> in the <a href="#">ONTAP 9 Documentation Center</a> to configure storage for the HA pair.

Precheck	Description
Data LIF Status Check	Checks if any of the nodes being replaced have non-local data LIFs. The nodes should not contain any data LIFs for which they are not the home owner. If one of the nodes contains non-local data LIFs, the task fails.
Cluster LIF Status	Checks whether the cluster LIFs are up for both nodes. If the cluster LIFs are down, the task fails.
ASUP Status Check	If ASUP notifications are not configured, the task fails. You must enable ASUP before beginning the controller replacement procedure.
CPU Utilization Check	Checks if the CPU utilization is more than 50% for any of the nodes being replaced. If the CPU usage is more than 50% for a considerable period of time, the task fails.
Aggregate Reconstruction Check	Checks if reconstruction is occurring on any data aggregates. If aggregate reconstruction is in progress, the task fails.
Node Affinity Job Check	Checks if any node affinity jobs are running. If node affinity jobs are running, the check fails.

3. After the controller replacement operation is started and the prechecks are completed the operation pauses, allowing you to collect output information that you might need later when configuring node3.
4. Run the below set of commands as directed by the controller replacement procedure on the system console.

You must run and save the output of the following commands individually:

- `vserver services name-service dns show`
- `network interface show -curr-node nodename -role cluster,intercluster,node-mgmt,cluster-mgmt, data`
- `network port show -node nodename -type physical`
- `service-processor show -node * -instance`
- `network fcp adapter show -node node_name`
- `network port ifgrp show`
- `network port vlan show`
- `system node show -instance -node nodename`
- `run -node node_name sysconfig`
- `storage aggregate show -node nodename`

- `volume show -node node_name`
- `network interface failover-groups show`
- `storage array config show -switch switch_name`
- `system license show -owner node_name`
- `storage encryption disk show`

## Correcting aggregate ownership if an ARL precheck fails

If the Aggregate Status Check fails, you must return aggregates owned by the partner node to the home owner node and initiate the precheck process again.

### Steps

1. Return the aggregates currently owned by the partner node to the home owner node:
 

```
storage aggregate relocation start -node source_node -destination destination-node -aggregate-list *
```
2. Verify that neither node1 nor node2 still owns aggregates for which it is the current owner (but not the home owner):

```
storage aggregate show -nodes node_name -is-home false -fields owner-name,home-name,state
```

### Example

The following example shows the output of the command when a node is both the current owner and home owner of aggregates:

```
cluster::> storage aggregate show -nodes node1 -is-home true -fields owner-name,home-name,state

aggregate    home-name    owner-name    state
-----
aggr1        node1        node1         online
aggr2        node1        node1         online
aggr3        node1        node1         online
aggr4        node1        node1         online

4 entries were displayed.
```

### After you finish

You must restart the controller replacement process:

```
system controller replace start -nodes node_names
```

## Licensing

When you set up a cluster, the setup wizard prompts you to enter the cluster-base license key. However, some features require additional licenses, which are issued as *packages* that include one or more features. Each node in the cluster must have its own key for each feature to be used in the cluster.

If you do not have new license keys, currently licensed features in the cluster are available to the new controller. However, using unlicensed features on the controller might put you out of compliance with your license agreement, so you should install the new license key or keys for the new controller after the upgrade is complete.

You can obtain new 2-character license keys for ONTAP 9.5 on the NetApp Support Site at [mysupport.netapp.com](http://mysupport.netapp.com). The keys are available in the *My Support* section under *Software licenses*. If the site does not have the license keys you need, you can contact your NetApp sales representative.

For detailed information about licensing, see the *System Administration Reference* in the [ONTAP 9 Documentation Center](#).

## Getting an IP address of an external key management server for storage encryption

After upgrading, you must immediately configure Storage Encryption and establish a cluster-wide authentication key to replace the previous node-level authentication keys.

### Steps

1. Install the necessary client and server SSL certificates required to communicate with key management servers by using the `security certificate install` command.
2. Configure Storage Encryption on all nodes by using the `security key-manager setup` command on each node.
3. Add the IP address for each key management server by using the `security key-manager add` command.
4. Verify that the same key management servers are configured and available on all nodes in the cluster:
 

```
security key-manager show -status
```
5. Create a new cluster-wide authentication key:
 

```
security key-manager create-key
```
6. Make a note of the new authentication key ID.
7. Rekey all SEDs with the new authentication key:
 

```
storage encryption disk modify -disk * -data-key-id authentication_key_id
```

## Managing authentication using KMIP servers

With ONTAP 9.5 and later, you can use KMIP servers to manage authentication keys.

### Steps

1. Add a new controller:
 

```
security key-manager setup -node new_controller_name
```
2. Add the key manager:
 

```
security key-manager -add key_management_server_ip_address
```
3. Verify that the key management servers are configured and available to all nodes in the cluster:
 

```
security key-manager show -status
```
4. Restore the authentication keys from all linked key management servers to the new node:
 

```
security key-manager restore -node new_controller_name
```
5. Rekey all self-encrypting disks with the new authentication key:
 

```
storage encryption disk modify -disk * -data-key-id nonMSID AK
```
6. If you use the Federal Information Processing Standard (FIPS), rekey all self-encrypting disks with the new authentication key:

```
storage encryption disk* modify -disk * -fips-key-id nonMSID AK
```

## Managing authentication using an onboard key manager

You can use an onboard key manager to manage authentication keys. If you plan to use an onboard key manager, you must record the passphrase and backup material before the beginning the upgrade.

### Steps

1. Verify the key management servers are available to all nodes in the cluster using the `security key-manager key show` command.

2. Rekey all self-encrypting disks with the new authentication key:

```
storage encryption disk modify -disk * -data-key-id nonMSID AK
```

3. If you use the Federal Information Processing Standard (FIPS), rekey all self-encrypting disks with the new authentication key:

```
storage encryption disk* modify -disk * -fips-key-id nonMSID AK
```

## Quiescing the SnapMirror relationships (optional)

Before you proceed with the replacement steps, you must ensure that all the SnapMirror relationships are quiesced. When a SnapMirror relationship is quiesced, it remains quiesced across reboots and failovers.

### Steps

1. Verify the SnapMirror relationship state on the destination cluster:

```
snapmirror show
```

**Note:** If the state is `Transferring`, you must abort those transfers by using the

```
snapmirror abort -destination vserver
```

command. The abort fails if the SnapMirror relationship is not in the `Transferring` state.

2. Quiesce all relationships between the cluster:

```
snapmirror quiesce -destination vserver
```

## Stage 2 Relocating and retiring node1

---

During Stage 2, you relocate the node1 aggregates and LIFs to node2, record node1 information, and then retire node1. This process is largely automated; the operation pauses to allow you to check its status. You must manually resume the operation.

### Steps

1. [Relocating non-root aggregates and NAS data LIFs owned by node1 to node2](#) on page 19
2. [Retiring node1](#) on page 21
3. [Preparing for netboot](#) on page 22

## Relocating non-root aggregates and NAS data LIFs owned by node1 to node2

Before you can replace node1 with node3, you must move the non-root aggregates and NAS data LIFs from node1 to node2 before eventually moving node1's resources to node3.

### Before you begin

The operation should already be paused when you begin the task; you must manually resume the operation.

### About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. Moving SAN LIFs is not necessary for cluster or service health during the upgrade. You must verify that the LIFs are healthy and located on appropriate ports after you bring node3 online.

**Note:** The home-owner for the aggregates and LIFs are not modified; only the current-owner is modified.

### Steps

1. Resume the aggregate relocation and NAS data LIF move operations:
 

```
system controller replace resume
```

All the non-root aggregates and NAS data LIFs are migrated from node1 to node2.

The operation pauses to allow you to verify whether all node1 non-root aggregates and non-SAN data lifs have been migrated to node2.
2. Check the status of the aggregate relocation and NAS data LIF move operations:
 

```
system controller replace show-details
```
3. With the operation still paused, verify that all the non-root aggregates are online and their state on node2 by entering the following command:
 

```
storage aggregate show -node node2 -state online -root false
```

### Example

The following example shows that the non-root aggregates on node2 are online:

```
cluster::> storage aggregate show -node node2 state online -root false
```

Aggregate	Size	Available	Used%	State	#Vols	Nodes	RAID Status
aggr_1	744.9GB	744.8GB	0%	online	5	node2	raid_dp, normal
aggr_2	825.0GB	825.0GB	0%	online	1	node2	raid_dp, normal

2 entries were displayed.

If the aggregates have gone offline or become foreign on node2, bring them online by using the following command on node2, once for each aggregate:

```
storage aggregate online -aggregate aggr_name
```

4. Verify that all the volumes are online on node2 by entering the following command on node2 and examining its output:

```
volume show -node node2 -state offline
```

If any volumes are offline on node2, bring them online by using the following command on node2, once for each volume:

```
volume online -vserver vservice-name -volume volume-name
```

The *vservice-name* to use with this command is found in the output of the previous `volume show` command.

5. If the ports currently hosting data LIFs will not exist on the new hardware, remove them from the broadcast domain:

```
network port broadcast-domain remove-ports
```

6. If any LIFs are down, set the administrative status of the LIFs to “up” by entering the following command, once for each LIF:

```
network interface modify -vserver vservice_name -lif LIF_name -home-node nodename -status-admin up
```

7. If you have interface groups or VLANs configured, complete the following substeps:

- a. If you have not already saved them, record VLAN and ifgrp information so you can recreate the VLANs and ifgrps on node3 after node3 is booted up.

- b. Remove the VLANs from the interface groups by entering the following command:

```
network port vlan delete -node nodename -port ifgrp -vlan-id VLAN_ID
```

**Note:** Follow the corrective action to resolve any errors that are suggested by the `vlan delete` command.

- c. Enter the following command and examine its output to see if there are any interface groups configured on the node:

```
network port ifgrp show -node nodename -ifgrp ifgrp_name -instance
```

### Example

The system displays interface group information for the node as shown in the following example:

```
cluster::> network port ifgrp show -node node1 -ifgrp a0a -instance
Node: node1
Interface Group Name: a0a
Distribution Function: ip
Create Policy: multimode_lacp
MAC Address: 02:a0:98:17:dc:d4
Port Participation: partial
Network Ports: e2c, e2d
Up Ports: e2c
Down Ports: e2d
```

- d. If any interface groups are configured on the node, record the names of those groups and the ports assigned to them, and then delete the ports by entering the following command, once for each port:

```
network port ifgrp remove-port -node nodename -ifgrp ifgrp_name -port
netport
```

## Relocating failed or vetoed aggregates

If any aggregates fail to relocate or are vetoed, you must take manually relocate the aggregates, or override either the vetoes or destination checks, if necessary.

### About this task

The relocation operation will have paused due to the error.

### Steps

1. Check the EMS logs to determine why the aggregate failed to relocate or was vetoed.

2. Relocate any failed or vetoed aggregates:

```
storage aggregate relocation start -node node1 -destination node2
aggregate-list * -ndocontroller-upgrade true
```

3. When prompted, enter y.

4. You can force relocation by using one of the following methods:

Option	Description
Overriding veto checks	Enter the following:  <pre>storage aggregate relocation start -override- vetoes * -ndocontroller-upgrade true</pre>
Overriding destination checks	Enter the following:  <pre>storage aggregate relocation start - overridedestination- checks * -ndo- controllerupgrade true</pre>

## Retiring node1

To retire node1, you resume the automated operation to disable the HA pair with node2 and shut node1 down properly. You must later remove node1 from the rack or chassis.

### Steps

1. Resume the operation:

```
system controller replace resume
```

2. Verify that node1 has been halted:

```
system controller replace show-details
```

### After you finish

You can decommission node1 after the upgrade is completed. See *Decommissioning the old system*.

## Preparing for netboot

After you install the new nodes later in the procedure, you might need to netboot them. The term *netboot* means you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you must put a copy of the ONTAP 9 boot image onto a web server that the system can access.

### Before you begin

- You must have verified that you can access an HTTP server with the system before and after installing node3 and node4.
- You must have verified that you have access to the NetApp Support Site at [mysupport.netapp.com](https://mysupport.netapp.com).  
You need to download the necessary system files for your platform and the correct version of ONTAP from the NetApp Support Site.

### About this task

You must netboot the new controllers if they do not have the same or later version of ONTAP 9 installed on them that is installed on the original controllers. After you install each new controller, you boot the system from the ONTAP 9 image stored on the web server. You can then download the correct files to the boot media device for subsequent system boots.

However, you do not need to netboot the controllers if the same or later version of ONTAP 9 is installed on them that is installed on the original controllers. If so, you can skip this section and proceed to [Stage 3 Installing and booting node3](#) on page 23.

### Steps

1. Download and extract the files used for performing the netboot of the system:
  - a. Download the appropriate `netboot.tgz` file from the software download section of the NetApp Support Site to a web-accessible directory.
  - b. Change to the web-accessible directory.
  - c. Extract the contents of the `netboot.tgz` file to the target directory:

```
tar -zxvf netboot.tgz
```

**Note:** If you are extracting the contents using Windows, do not use winzip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

Your directory listing should contain the following directory:

```
netboot/
```

2. Download the `image.tgz` file from the NetApp Support Site to the web-accessible directory that you chose in Step 1a.

Your directory listing should contain the following files and directory:

```
image.tgz
netboot.tgz
netboot/
```

You will use information in the directories in Stage 3.

## Stage 3 Installing and booting node3

---

During Stage 3, you install and boot node3, map the cluster and node-management ports from node1 to node3, and verify the node3 installation. Non-SAN data LIFs and non-root aggregates belonging to node1 are moved from node2 to node3.

### About this task

#### Steps

1. [Installing and booting node3](#) on page 23
2. [Setting the FC or UTA/UTA2 configuration on node3](#) on page 27
3. [Mapping ports from node1 to node3](#) on page 34
4. [Joining the quorum when a node has a different set of network ports](#) on page 37
5. [Verifying the node3 installation](#) on page 39
6. [Moving non-root aggregates and NAS data LIFs owned by node1 from node2 to node3](#) on page 40

## Installing and booting node3

You need to install node3 in the rack, transfer node1's connections to node3, boot node3, and install ONTAP. You then need to reassign any of node1's spare disks, any disks belonging to the root volume, and any non-root aggregates not relocated to node2 earlier.

### About this task

The relocation operation is paused at the beginning of this stage. This process is largely automated; the operation pauses to allow you to check its status. You must manually resume the operation. In addition, you must verify the SAN LIFs have successfully moved to node3.

You need to netboot node3 if it does not have the same or a later version of ONTAP 9 that is installed on node1. After you install node3, boot it from the ONTAP 9 image stored on the web server. You can then download the correct files to the boot media device for subsequent system boots.

**Important:** If you are upgrading a V-Series system connected to storage arrays or a system with FlexArray Virtualization software that is connected to storage arrays, you need to complete Step [1](#) on page 23 through Step [18](#) on page 26, then leave this section and follow instructions in the [Configuring FC ports on node3](#) on page 27 and [Checking and configuring UTA/UTA2 ports on node3](#) on page 29 sections as needed, entering commands in Maintenance mode. You must then return to this section and resume with Step [20](#) on page 26.

However, if you are upgrading a system with storage disks, you need to complete this entire section and then go to the [Configuring FC ports on node3](#) on page 27 and [Checking and configuring UTA/UTA2 ports on node3](#) on page 29 sections, entering commands at the cluster prompt.

#### Steps

1. Make sure that you have rack space for node3.

If node1 and node2 were in separate chassis, you can put node3 in the same rack location as node1. However, if node1 was in the same chassis with node2, then you need to put node3 into its own rack space, preferably close to the location of node1.

2. Install node3 in the rack, following instructions in the *Installation and Setup Instructions* for your node model.

**Note:** If you are upgrading to a system with both nodes in the same chassis, install node4 in the chassis as well as node3. If you do not, when you boot node3, the node will behave as if it were in a dual-chassis configuration, and when you boot node4, the interconnect between the nodes will not come up.

3. Cable node3, moving the connections from node1 to node3.

Cable the following connections, using the instructions in the *Installation and Setup Instructions* or the *FlexArray Virtualization Installation Requirements and Reference* for the node3 platform, the appropriate disk shelf guide, and the *ONTAP 9 High-Availability Configuration Guide*:

- Console (remote management port)
- Cluster ports
- Data ports
- Cluster and node management ports
- Storage
- SAN configurations: iSCSI Ethernet and FC switch ports

**Note:** You do not need to move the interconnect card or the cluster interconnect cable connection from node1 to node3 because most platform models have a unique interconnect card model.

4. Turn on the power to node3, and then interrupt the boot process by pressing Ctrl-C at the console terminal to access the boot environment prompt.

If you are upgrading to a system with both nodes in the same chassis, node4 also reboots. However, you can disregard the node4 boot until later.

**Note:** When you boot node3, you might see the following warning message:

```
WARNING: The battery is unfit to retain data during a power
         outage. This is likely because the battery is
         discharged but could be due to other temporary
         conditions.
         When the battery is ready, the boot process will
         complete and services will be engaged.
         To override this delay, press 'c' followed by 'Enter'
```

5. If you see the warning message in Step 4 on page 24, take the following actions:
  - a. Check for any console messages that might indicate a problem other than a low NVRAM battery, and, if necessary, take any required corrective action.
  - b. Allow the battery to charge and the boot process to complete.

**Attention:** Do not override the delay; failure to allow the battery to charge could result in a loss of data.

**Note:** Refer to [Preparing for netboot](#) on page 22

6. Configure the netboot connection by choosing one of the following actions.

**Note:** You should use the management port and IP as the netboot connection. Do not use a data LIF IP or a data outage might occur while the upgrade is being performed.

If Dynamic Host Configuration Protocol (DHCP) is...	Then...
Running	Configure the connection automatically by entering the following command at the boot environment prompt:  <b>ifconfig e0M -auto</b>
Not running	Manually configure the connection by entering the following command at the boot environment prompt:  <b>ifconfig e0M -addr=filer_addr -mask=netmask -gw=gateway -dns=dns_addr domain=dns_domain</b>  <i>filer_addr</i> is the IP address of the storage system.  <i>netmask</i> is the network mask of the storage system.  <i>gateway</i> is the gateway for the storage system.  <i>dns_addr</i> is the IP address of a name server on your network.  <i>dns_domain</i> is the Domain Name Service (DNS) domain name. If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL; you need only the server's host name.  <b>Note:</b> Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.

7. Perform netboot by entering the following command on node3:

```
netboot http://path_to_the_web-accessible_directory/netboot/kernel
```

The *path\_to\_the\_web-accessible\_directory* should lead to where you downloaded the `netboot.tgz` in Step 1a in the section [Preparing for netboot](#) on page 22.

**Note:** Do not interrupt the boot.

8. From the boot menu, select option (7) Install new software first.

This menu option downloads and installs the new ONTAP image to the boot device.

**Note:** Disregard the following message: "This procedure is not supported for Non-Disruptive Upgrade on an HA pair". The note applies to nondisruptive upgrades of ONTAP, and not upgrades of controllers.

9. If you are prompted to continue the procedure, enter **y**, and when prompted for the package, enter the URL `http://path_to_the_web-accessible_directory/image.tgz`.
10. Complete the following substeps:

- a. Enter **n** to skip the backup recovery when you see the following prompt:

```
Do you want to restore the backup configuration now? {y|n}
n
```

- b. Reboot by entering **y** when you see the following prompt:

```
The node must be rebooted to start using the newly installed
software. Do you want to reboot now? {y|n}
```

The controller module reboots but stops at the boot menu because the boot device was reformatted and the configuration data needs to be restored.

11. Select reboot option 8 from the boot menu and interrupt the autoboot by pressing Ctrl-C.
12. On node2, check the system date, time, and time zone:  
`date`
13. On node3, check the date by entering the following command at the boot environment prompt:  
`show date`
14. If necessary, set the date on node3:  
`set date mm/dd/yyyy`
15. On node3, check the time by entering the following command at the boot environment prompt:  
`show time`
16. If necessary, set the time on node3:  
`set time hh:mm:ss`
17. If necessary, set the partner system ID on node3:  
`setenv partner-sysid node2_sysid`
  - a. Save the settings:  
`saveenv`
18. On the new node, in boot loader, the `partner-sysid` parameter needs to be set. For node3, `partner-sysid` needs to be that of node2. Verify the `partner-sysid` for node3 using the following command:

```
printenv partner-sysid
```

19. Take one of the following actions:

If your system...	Description
Has disks and no back-end storage	Go to Step <a href="#">20</a> on page 26.
Is a V-Series system or a system with FlexArray Virtualization software connected to storage arrays	<ol style="list-style-type: none"> <li>a. Go to <a href="#">Setting the FC or UTA/UTA2 configuration on node3</a> on page 27 and complete the subsections in it.</li> <li>b. Return to this section and complete the remaining steps, beginning with Step <a href="#">20</a> on page 26.</li> </ol> <p><b>Important:</b> You must reconfigure FC onboard ports, CNA onboard ports, and CNA cards before you boot ONTAP on the V-Series or system with FlexArray Virtualization software.</p>

20. Add the FC initiator ports of the new node to the switch zones.  
If your system has a tape SAN, then you need zoning for the initiators. If required, modify the onboard ports to initiator by referring to the [Configuring FC ports on node3](#) on page 27. See your storage array and zoning documentation for further instructions on zoning.
21. Add the FC initiator ports to the storage array as new hosts, mapping the array LUNs to the new hosts.  
See your storage array and zoning documentation for instructions.
22. Modify the World Wide Port Name (WWPN) values in the host or volume groups associated with array LUNs on the storage array.  
Installing a new controller module changes the WWPN values associated with each onboard FC port.

23. If your configuration uses switch-based zoning, adjust the zoning to reflect the new WWPN values.
24. If NetApp Storage Encryption (NSE) is in use on this configuration, the `setenv bootarg.storageencryption.support` command must be set to **true**, and the `kmip.init.maxwait` variable needs to be set to **off** to avoid a boot loop after the node1 configuration is loaded:
 

```
setenv bootarg.storageencryption.support true
setenv kmip.init.maxwait off
```
25. Boot node into boot menu by running `boot_ontap menu`.  
If you do not have FC or UTA/UTA2 configuration, execute Step 15 on page 32 so that node3 can recognize node1's disks.
26. For V-Series systems and systems with FlexArray Virtualization software connected to storage arrays proceed to Step 15 on page 32.

## Setting the FC or UTA/UTA2 configuration on node3

If node3 has onboard FC ports, onboard unified target adapter (UTA/UTA2) ports, or a UTA/UTA2 card, you must configure the settings before completing the rest of the procedure.

### About this task

You might need to complete the section [Configuring FC ports on node3](#) on page 27, the section [Checking and configuring UTA/UTA2 ports on node3](#) on page 29, or both sections.

**Note:** NetApp marketing materials might use the term *UTA2* to refer to CNA adapters and ports. However, the CLI uses the term *CNA*.

If node3 does not have onboard FC ports, onboard UTA/UTA2 ports, or a UTA/UTA2 card, and you are upgrading a system with storage disks, you can skip to the section [Mapping ports from node1 to node3](#) on page 34. However, if you have a V-Series system or a system with FlexArray Virtualization software with storage arrays, and node3 does not have onboard FC ports, onboard UTA/UTA ports, or a UTA/UTA2 card, return to the section [Installing and booting node3](#) on page 23 and resume the section at Step 20 on page 26.

### Choices

- [Configuring FC ports on node3](#) on page 27
- [Checking and configuring UTA/UTA2 ports on node3](#) on page 29

## Configuring FC ports on node3

If node3 has FC ports, either onboard or on an FC adapter, you must set port configurations on the node before you bring it into service because the ports are not preconfigured. If the ports are not configured, you might experience a disruption in service.

### Before you begin

You must have the values of the FC port settings from node1 that you saved in the section [Preparing the nodes for upgrade](#) on page 13.

### About this task

You can skip this section if your system does not have FC configurations. If your system has onboard UTA/UTA2 ports or a UTA/UTA2 card, you configure them in [Checking and configuring UTA/UTA2 ports on node3](#) on page 29.

**Important:** If your system has storage disks, enter the commands in this section at the cluster prompt. If you have a 'V-Series system' or have FlexArray Virtualization Software and are connected to storage arrays, enter commands in this section in Maintenance mode.

### Steps

1. Compare the FC settings on node3 with the settings that you captured earlier from node1.
2. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	<p>In maintenance mode (option 5 at boot menu), modify the FC ports on node3 as needed by entering one of the following commands:</p> <ul style="list-style-type: none"> <li>• To program target ports: <b>system node hardware unified-connect modify -type   -t target -adapter port_name</b></li> <li>• To program initiator ports: <b>system node hardware unified-connect modify -type   -t initiator -adapter port_name</b></li> </ul> <p>-type is the FC4 type: target or initiator.</p>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<p>In maintenance mode (option 5 at boot menu), modify the FC ports on node3 as needed by entering the following command: <b>ucadmin modify -m fc -t initiator -f adapter_port_name</b></p> <p>-t is the FC4 type, target or initiator.</p> <p><b>Note:</b> The FC ports must be programmed as initiators.</p>

3. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	<p>Verify the new settings by entering the following command and examining the output:</p> <p><b>system node hardware unified-connect show</b></p>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<p>Verify the new settings by entering the following command and examining the output:</p> <p><b>ucadmin show</b></p>

4. Exit Maintenance mode by entering the following command:
 

```
halt
```
5. After you enter the command, wait until the system stops at the boot environment prompt.
6. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	<ul style="list-style-type: none"> <li>If node3 has a UTA/UTA2 card or UTA/UTA2 onboard ports, go to the section <a href="#">Checking and configuring UTA/UTA2 ports on node3</a> on page 29.</li> <li>If node3 does not have a UTA/UTA2 card or UTA/UTA2 onboard ports, skip the section <a href="#">Checking and configuring UTA/UTA2 ports on node3</a> on page 29 and go to <a href="#">Mapping ports from node1 to node3</a> on page 34.</li> </ul>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<ul style="list-style-type: none"> <li>If node3 has a UTA/UTA2 card or UTA/UTA2 onboard ports, go to the section <a href="#">Checking and configuring UTA/UTA2 ports on node3</a> on page 29.</li> <li>If node3 does not have a UTA/UTA2 card or UTA/UTA2 onboard ports, skip the section <a href="#">Checking and configuring UTA/UTA2 ports on node3</a> on page 29 and return to <a href="#">Installing and booting node3</a> on page 23 and resume the section at <a href="#">20</a> on page 26.</li> </ul>

## Checking and configuring UTA/UTA2 ports on node3

If node3 has onboard UTA/UTA2 ports or a UTA/UTA2 card, you must check the configuration of the ports and possibly reconfigure them, depending on how you want to use the upgraded system.

### Before you begin

You must have the correct SFP+ modules for the UTA/UTA2 ports.

### About this task

If you want to use a Unified Target Adapter (UTA/UTA2) port for FC, you must first verify how the port is configured.

**Note:** NetApp marketing materials might use the term UTA2 to refer to CNA adapters and ports. However, the CLI uses the term CNA.

You can use the `ucadmin show` command to verify the current port configuration:

```

AFFRTPSE::> system node hardware unified-connect

```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
node1	0e	fc	target	-	-	online
node1	0f	fc	target	-	-	online
node1	0g	cna	target	-	-	online
node1	0h	cna	target	-	-	online
node1	0e	fc	target	-	-	online
node1	0f	fc	target	-	-	online
node1	0g	cna	target	-	-	online
node1	0h	cna	target	-	-	online

```
8 entries were displayed.
```

UTA/UTA2 ports can be configured into native FC mode or UTA/UTA2 mode. FC mode supports FC initiator and FC target; UTA/UTA2 mode allows concurrent NIC and FCoE traffic sharing the same 10GbE SFP+ interface and supports FC targets.

UTA/UTA2 ports might be found on an adapter or on the controller, and have the following configurations, but you should check the configuration of the UTA/UTA2 ports on the node3 and change it, if necessary:

- UTA/UTA2 cards ordered when the controller is ordered are configured before shipment to have the personality you request.
- UTA/UTA2 cards ordered separately from the controller are shipped with the default FC target personality.
- Onboard UTA/UTA2 ports on new controllers are configured before shipment to have the personality you request.

**Attention:** If your system has storage disks, you enter the commands in this section at the cluster prompt unless directed to enter Maintenance mode. If you have a V-Series system or have FlexArray Virtualization Software and are connected to storage arrays, you enter commands in this section at the Maintenance mode prompt. You must be in Maintenance mode to configure UTA/UTA2 ports.

## Steps

1. Check how the ports are currently configured by entering one of the following commands on node3:

If the system...	Then...
Has storage disks	<code>system node hardware unified-connect show</code>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<code>ucadmin show</code>

The system displays output similar to the following examples:

```
cluster1::> system node hardware unified-connect show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
f-a	0e	fc	initiator	-	-	online
f-a	0f	fc	initiator	-	-	online
f-a	0g	cna	target	-	-	online
f-a	0h	cna	target	-	-	online
f-b	0e	fc	initiator	-	-	online
f-b	0f	fc	initiator	-	-	online
f-b	0g	cna	target	-	-	online
f-b	0h	cna	target	-	-	online

```
12 entries were displayed.
```

```
*> ucadmin show
```

Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
0e	fc	initiator	-	-	online
0f	fc	initiator	-	-	online
0g	cna	target	-	-	online

```

0h      cna      target    -      -      online
0e      fc       initiator -      -      online
0f      fc       initiator -      -      online
0g      cna      target    -      -      online
0h      cna      target    -      -      online
*>

```

2. If the current SFP+ module does not match the desired use, replace it with the correct SFP+ module.

Contact your NetApp representative to obtain the correct SFP+ module.

3. Examine the output of the `system node hardware unified-connect show` or `ucadmin show` command and determine whether the UTA/UTA2 ports have the personality you want.
4. Take one of the following actions:

If the UTA/UTA2 ports...	Then...
Do not have the personality that you want	Go to Step 5.
Have the personality that you want	Skip Step 5 through Step 12 and go Step 13.

5. Take one of the following actions:

If you are configuring...	Then...
Ports on a UTA/UTA2 card	Go to Step 7.
Onboard UTA/UTA2 ports	Skip Step 7 and go to Step 8.

6. If the adapter is in initiator mode, and if the UTA/UTA2 port is online, take the UTA/UTA2 port offline:

**storage disable adapter *adapter\_name***

Adapters in target mode are automatically offline in Maintenance mode.

7. If the current configuration does not match the desired use, change the configuration as needed:

**ucadmin modify -m *fc|cna* -t *initiator|target* *adapter\_name***

- `-m` is the personality mode, `fc` or `cna`.
- `-t-t` is the FC4 type, `target` or `initiator`.

**Note:** You need to use FC initiator for tape drives, FlexArray Virtualization systems, and MetroCluster configurations. You need to use the FC target for SAN clients.

8. Verify the settings:

**ucadmin show**

9. Verify the settings:

If the system...	Then...
Has storage disks	<b>system node hardware unified-connect show</b>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<b>ucadmin show</b>

**Example**

The output in the following examples shows that the FC4 type of adapter 1b is changing to initiator and that the mode of adapters 2a and 2b is changing to cna:

```
cluster1::> system node hardware unified-connect show
Node   Adapter  Current Mode   Current Type   Pending Mode   Pending Type   Admin Status
-----
f-a    1a       fc      initiator -       -       initiator online
f-a    1b       fc      target  -       -       initiator online
f-a    2a       fc      target  cna     -       -       online
f-a    2b       fc      target  cna     -       -       online
4 entries were displayed.
```

```
*> ucadmin show
Adapter  Current Mode   Current Type   Pending Mode   Pending Type   Admin Status
-----
1a       fc      initiator -       -       initiator online
1b       fc      target  -       -       initiator online
2a       fc      target  cna     -       -       online
2b       fc      target  cna     -       -       online
*>
```

10. Place any target ports online by entering one of the following commands, once for each port:

If the system...	Then...
Has storage disks	<code>network fcp adapter modify -node node_name - adapter adapter_name -state up</code>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<code>fcp config adapter_name up</code>

11. Cable the port.

12. Take one of the following actions:

If the system...	Then...
Has storage disks.	Go to <a href="#">Mapping ports from node1 to node3</a> on page 34.
Is a V-series system or has FlexArray Virtualization Software and is connected to storage arrays	Return to the <a href="#">Stage 3 Installing and booting node3</a> on page 23 and resume the section at <a href="#">20</a> on page 26

13. Exit Maintenance mode by entering the following command:

```
halt
```

14. Boot node into boot menu by running `boot_ontap` menu.

15. On node3, go to the boot menu and using 22/7, select the hidden option `boot_after_controller_replacement`. At the prompt, enter node1 to reassign the disks of node1 to node3, as per the following example.

**Example**

```
LOADER-A> boot_ontap menu
...
*****
*
* Press Ctrl-C for Boot Menu. *
```

```

*
*****
.
.
Please choose one of the following:

(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
Selection (1-9)? 22/7
.
.
.
(boot_after_controller_replacement) Boot after controller upgrade
(9a) Unpartition all disks and remove their ownership information.
(9b) Clean configuration and initialize node with partitioned disks.
(9c) Clean configuration and initialize node with whole disks.
(9d) Reboot the node.
(9e) Return to main boot menu.

Please choose one of the following:

(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
Selection (1-9)? boot_after_controller_replacement
.
This will replace all flash-based configuration with the last backup to
disks. Are you sure you want to continue?: yes
.
.
Controller Replacement: Provide name of the node you would like to replace: <name of the node being
replaced>
.
.
Changing sysid of node <node being replaced> disks.
Fetched sanown old_owner_sysid = 536953334 and calculated old sys id = 536953334
Partner sysid = 4294967295, owner sysid = 536953334
.
.
.
Terminated
<node reboots>
.
.
System rebooting...
.
Restoring env file from boot media...
copy_env_file:scenario = head upgrade
Successfully restored env file from boot media...
.
.
System rebooting...
.
.
WARNING: System ID mismatch. This usually occurs when replacing a boot device or NVRAM cards!
Override system ID? {y|n} y

Login:
...

```

16. If the system goes into a reboot loop with the message `no disks found`, this is because it has reset the ports back to the target mode and therefore is unable to see any disks. Continue with Steps [17](#) on page 33 to [22](#) on page 33 to resolve this.
17. Press `Ctrl-C` and break the loop.
18. In loader prompt, use the command `boot_ontap` menu and enter maintenance mode.
19. In maintenance mode, use the command `ucadmin show` to display all the previously set initiator ports that are now in target mode. Change the ports back to initiator mode using the command `ucadmin modify -m fc -t initiator -f <adapter name>`.
20. Enter `ucadmin show` to verify that the ports have been changed to initiator mode.
21. Exit maintenance mode by entering the command `halt`.
22. In loader prompt, use the command `boot_ontap`.

Now, on booting, the node is able to see all the disks that were previously assigned to it and can boot up as expected.

## Mapping ports from node1 to node3

You must verify that the physical ports on node1 map correctly to the physical ports on node3, which will allow node3 to communicate with other nodes in the cluster and with the network after the upgrade.

### About this task

Capture information about the ports on the new nodes from the *Hardware Universe* at [hww.netapp.com](http://hww.netapp.com). You will use the information later in this section.

Port settings might vary, depending on the model of the nodes. You must make the port and LIF configuration on the original node compatible with the planned use and configuration of the new node. This is because the new node replays the same configuration when it boots, which means that when you boot node3, ONTAP will try to host LIFs on the same ports that were used on node1.

Therefore, if the physical ports on node1 do not map directly to the physical ports on node3, then software configuration changes will be required to restore cluster, management, and network connectivity after the boot. In addition, if the cluster ports on node1 do not directly map to the cluster ports on node3, node3 might not automatically rejoin quorum when it is rebooted until you change the software configuration to host the cluster LIFs on the correct physical ports.

### Steps

1. Record all the node1 cabling information for node1, the ports, broadcast domains, and IPspaces, in the table:

LIF	Node 1 ports	Node1 IPspaces	Node1 broadcast domains	Node3 ports	Node3 IPspaces	Node3 broadcast domains
Cluster 1						
Cluster 2						
Cluster 3						
Cluster 4						
Node management						
Cluster management						
Data 1						
Data 2						
Data 3						
Data 4						
SAN						
Intercluster port						

2. Record all the cabling information for node3, the ports, broadcast domains, and IPspaces in the table.
3. Follow these steps to verify if the setup is a two-node switchless cluster:

- a. Set the privilege level to advanced:

```
cluster::> set -privilege advanced
```

- b. Verify if the setup is a two-node switchless cluster:

```
cluster::> network options switchless-cluster show
```

#### Example

```
cluster::*> network options switchless-cluster show
Enable Switchless Cluster: false/true
```

The value of this command output must match the physical state of the system.

- c. Return to the administration privilege level:

```
cluster::*> set -privilege admin
cluster::>
```

4. Follow these steps to place node3 into quorum:
  - a. Boot node3. See the [Installing and booting node3](#) on page 23 section to boot the node if you have not already done so.
  - b. Verify that the new cluster ports are in the Cluster broadcast domain:

```
network port show -node node -port port -fields broadcast-domain
```

#### Example

The following example shows that port e0a is in the Cluster domain on node3:

```
cluster::> network port show -node node3 -port e0a -fields
broadcast-domain

node      port broadcast-domain
-----
node3    e1a Cluster
```

- c. If the cluster ports are not in the Cluster broadcast-domain, add them with the following command:

```
broadcast-domain add-ports -ip-space Cluster -broadcast-domain Cluster
-ports node:port
```

#### Example

This example adds Cluster port e1b on node3:

```
network port modify -node node3 -port e1b -ip-space Cluster -mtu 9000
```

- d. Add the correct ports to the Cluster broadcast domain:

```
network port modify -node -port -ip-space Cluster -mtu 9000
```

**Example**

This example adds Cluster port e1b on node4:

```
network port modify -node node4 -port e1b -ipspace Cluster -mtu 9000
```

**Note:** For a MetroCluster configuration, you may not be able to change the broadcast domain of a port because it is associated with a port hosting the LIF of a sync-destination SVM and see errors similar to, but not restricted, to the following:

```
command failed: This operation is not permitted on a Vserver that
is
    configured as the destination of a MetroCluster Vserver
relationship.
```

Enter the following command from the corresponding sync-source SVM on the remote site to reallocate the sync-destination LIF to an appropriate port:

```
metrocluster vserver resync -vserver vserver_name
```

- e. Migrate the cluster LIFs to the new ports, once for each LIF:

```
network interface migrate -vserver Cluster -lif lif_name -source-node
node3 -destination-node node3 -destination-port port_name
```

- f. Modify the home port of the cluster LIFs:

```
network interface modify -vserver Cluster -lif lif_name -home-port
port_name
```

- g. Remove the old ports from the Cluster broadcast-domain:

```
network port broadcast-domain remove-ports
```

**Example**

The following command removes port e0d on node3:

```
network port broadcast-domain remove-ports -ipspace Cluster -
broadcast-domain Cluster -ports node3:e0d
```

- h. Verify that node3 has rejoined quorum:

```
cluster show -node node3 -fields health
```

5. Adjust the broadcast domains hosting your cluster LIFs and node-management/cluster-management LIFs. Ensure that each broadcast domain contains the correct ports. A port cannot be moved between broadcast domains if it is hosting or is home to a LIF, so you may need to migrate and modify the LIFs as follows:

- a. Display the home port of a LIF:

```
network interface show -fields home-node,home-port
```

- b. Display the broadcast domain containing this port:

```
network port broadcast-domain show -ports node_name:port_name
```

- c. Add or remove ports from broadcast domains:

```
network port broadcast-domain add-ports
network port broadcast-domain remove-ports
```

- d. Modify a LIF's home port:

```
network interface modify -vserver vserver -lif lif_name -home-port
port_name
```

6. Adjust the intercluster broadcast domains and migrate the intercluster LIFs, if necessary, using the same commands shown in Step 5.
7. Adjust any other broadcast domains and migrate the data LIFs, if necessary, using the same commands shown in Step 5.
8. If there were any ports on node1 that no longer exist on node3, follow these steps to delete them:
  - a. Access the advanced privilege level on either node:
 

```
set -privilege advanced
```
  - b. To delete the ports:
 

```
network port delete -node node_name -port port_name
```
  - c. Return to the admin level:
 

```
set -privilege admin
```
9. Adjust all the LIF failover groups:
 

```
network interface modify -failover-group failover_group -failover-policy failover_policy
```

#### Example

The following command sets the failover policy to `broadcast-domain-wide` and uses the ports in failover group `fg1` as failover targets for LIF `data1` on node3:

```
network interface modify -vserver node3 -lif data1 failover-policy broadcast-domain-wide -failover-group fg1
```

See the *ONTAP 9 Network Management Guide* or *ONTAP 9 Commands: Manual Page Reference* for more information.

10. Verify the changes on node3:
 

```
network port show -node node3
```

## Joining the quorum when a node has a different set of network ports

The node with the new controller boots and attempts to join the cluster automatically at first; however if the new node has a different set of network ports, you must perform the following steps to ensure that the node successfully joins the quorum.

#### About this task

You can use these instructions for any relevant node. Node3 is used throughout the following sample.

#### Steps

1. Verify that the new cluster ports are in the Cluster broadcast domain by entering the following command and checking its output:
 

```
network port show -node node -port port -fields broadcast-domain
```

#### Example

The following example shows that port `e1a` is in the Cluster domain on node3:

```
cluster::> network port show -node node3 -port ela -fields broadcast-
domain
node    port    broadcast-domain
-----  ----  -
node3   ela    Cluster
```

2. Add the correct ports to the Cluster broadcast domain by entering the following command and checking its output:

```
network port modify -node -port -ip-space Cluster -mtu 9000
```

#### Example

This example adds Cluster port e1b on node3:

```
network port modify -node node3 -port e1b -ip-space Cluster -mtu 9000
```

3. Migrate the cluster LIFs to the new ports, once for each LIF, using the following command:
 

```
network interface migrate -vserver Cluster -lif lif_name -source-node
node3 -destination-node node3 -destination-port port_name
```
4. Modify the home port of the cluster LIFs as follows:
 

```
network interface modify -vserver Cluster -lif lif_name -home-port
port_name
```
5. If the cluster ports are not in the Cluster broadcast-domain, add them with the following command:
 

```
network port broadcast-domain add-ports -ip-space Cluster -
broadcastdomain Cluster -ports node:port
```
6. Remove the old ports from the Cluster broadcast domain. You can use for any relevant node. The following command removes port e0d on node3:
 

```
network port broadcast-domain remove-ports network port broadcast-domain
remove-ports -ip-space Cluster -broadcast-domain Cluster -ports node3:e0d
```
7. Verify the node has rejoined quorum as follows:
 

```
cluster show -node node3 -fields health
```
8. Adjust the broadcast domains hosting your cluster LIFs and node-management/cluster management LIFs. Ensure that each broadcast domain contains the correct ports. A port cannot be moved between broadcast domains if it is hosting or is home to a LIF, so you might need to migrate and modify the LIFs as follows:
  - a. Display the home port of a LIF:
 

```
network interface show -fields home-node,home-port
```
  - b. Display the broadcast domain containing this port:
 

```
network port broadcast-domain show -ports node_name:port_name
```
  - c. Add or remove ports from broadcast domains:
 

```
network port broadcast-domain add-ports network port broadcast-domain
remove-port
```
  - d. Modify a home port of a LIF:
 

```
network interface modify -vserver vserver -lif lif_name -home-port
port_name
```

Adjust the intercluster broadcast domains and migrate the intercluster LIFs, if necessary. The data LIFs remain unchanged.

## Verifying the node3 installation

After you install and boot node3, you need to verify that it is installed correctly. You must wait for node3 to join quorum and then resume the relocation operation.

### About this task

At this point in the procedure, the operation will have paused as node3 joins quorum.

### Steps

1. Verify that node3 has joined quorum:  
`cluster show -node node3 -fields health`
2. Verify that node3 is part of the same cluster as node2 and that it is healthy by entering the following command:  
`cluster show`
3. Check the status of the operation and verify that the configuration information for node3 is the same as node1:

```
system controller replace show-details
```

If the configuration is different for node3, a system disruption might occur later in the procedure.

## Re-creating VLANs, ifgrps, and broadcast domains on node3

After you confirm that node3 is in quorum and can communicate with node2, you must re-create node1's VLANs, ifgrps, and broadcast domains on node3. You must also add the node3 ports to the newly re-created broadcast domains.

### About this task

For more information on creating and re-creating VLANs, ifgrps, and broadcast domains, see the *ONTAP 9.5 Administration Reference*.

### Steps

1. Re-create the VLANs on node3 using the node1 information recorded in the [Relocating non-root aggregates and NAS data LIFs owned by node1 to node2](#) on page 19:  
`network port vlan create -node node_name -vlan vlan-names`
2. Re-create the ifgrps on node3 using the node1 information recorded in the [Relocating non-root aggregates and NAS data LIFs owned by node1 to node2](#) on page 19 section:  
`network port ifgrp create -node node_name -ifgrp port_ifgrp_names-distr-func`
3. Re-create the broadcast domains on node3 using the node1 information recorded in the [Relocating non-root aggregates and NAS data LIFs owned by node1 to node2](#) on page 19:  
`network port broadcast-domain create -ip-space Default -broadcast-domain broadcast_domain_names -mtu mtu_size -ports node_name:port_name,node_name:port_name`
4. Add the node3 ports to the newly re-created broadcast domains:

```
network port broadcast-domain add-ports -broadcast-domain
broadcast_domain_names -ports node_name:port_name,node_name:port_name
```

## Moving non-root aggregates and NAS data LIFs owned by node1 from node2 to node3

After you verify the node3 installation and before you relocate aggregates from node2 to node3, you need to move the NAS data LIFs belonging to node1 that are currently on node2 from node2 to node3. You also need to verify that the SAN LIFs exist on node3.

### About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. Moving SAN LIFs is not necessary for cluster or service health during the upgrade. SAN LIFs are not moved unless they need to be mapped to new ports. You will verify that the LIFs are healthy and located on appropriate ports after you bring node3 online.

### Steps

1. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following tasks:

- Cluster quorum check
- System ID check
- Image version check
- Target platform check
- Network reachability check

The operation pauses at this stage in the network reachability check.

2. Manually verify that the network and all VLANs, ifgrps, and broadcast domains have been configured correctly.
3. Resume the relocation operation:

```
system controller replace resume
```

To complete the "Network Reachability" phase, ONTAP network configuration must be manually adjusted to match the new physical network configuration of the hardware. This includes assigning network ports to the correct broadcast domains, creating any required ifgrps and VLANs, and modifying the home-port parameter of network interfaces to the appropriate ports. Refer to the "Using aggregate relocation to upgrade controller hardware on a pair of nodes running ONTAP 9.x" documentation, Stages 3 and 5. Have all of these steps been manually completed? [y/n]

4. Enter y to continue.
5. The system performs the following checks:
  - Cluster health check
  - Disk status check
  - Cluster LIF status check

After performing the three checks, the system relocates the non-root aggregates and NAS data LIFs owned by node1 to the new controller, node3.

The system pauses once the resource relocation is complete.

6. Check the status of the aggregate relocation and NAS data LIF move operations:

```
system controller replace show-details
```

7. Verify that the non-root aggregates and NAS data LIFs have been successfully relocated to node3.

If any aggregates fail to relocate or are vetoed, you must manually relocate the aggregates, or override either the vetoes or destination checks, if necessary. See the section [Relocating failed or vetoed aggregates](#) on page 21 for more information.

8. Ensure that the SAN LIFs are on the correct ports on node3 by completing the following substeps:

- a. Enter the following command and examine its output:

```
network interface show -data-protocol iscsi|fcp -home-node node3
```

### Example

The system returns output similar to the following example:

```
cluster::> net int show -data-protocol iscsi|fcp -home-node node3
-----
Vserver Logical Status Network Current Current Is
Interface Admin/Oper Address/Mask Node Port Home
-----
vs0
  a0a      up/down  10.63.0.53/24  node3  a0a    true
  data1    up/up     10.63.0.50/18  node3  e0c    true
  rads1    up/up     10.63.0.51/18  node3  e1a    true
  rads2    up/down  10.63.0.52/24  node3  e1b    true
vs1
  lif1     up/up     172.17.176.120/24 node3  e0c    true
  lif2     up/up     172.17.176.121/24 node3  e1a    true
```

- b. If node3 has any SAN LIFs or groups of SAN LIFs that are on a port that did not exist on node1 or that need to be mapped to a different port, move them to an appropriate port on node3 by completing the following substeps:

- i. Set the LIF status to down by entering the following command:

```
network interface modify -vserver vserver_name -lif lif_name -
status-admin down
```

- ii. Remove the LIF from the port set:

```
portset remove -vserver vserver_name -portset portset_name -port-
name port_name
```

- iii. Enter one of the following commands:

- Move a single LIF by entering the following command:

```
network interface modify -vserver vserver_name -lif lif_name -
home-port new_home_port
```

- Move all the LIFs on a single nonexistent or incorrect port to a new port by entering the following command:

```
network interface modify {-home-port port_on_node1 -home-node
node1 -role data} -home-port new_home_port_on_node3
```

- Add the LIFs back to the port set:

```
portset add -vserver vserver_name -portset portset_name -port-
name port_name
```

**Note:** You need to ensure that you move SAN LIFs to a port that has the same link speed as the original port.

- c. Modify the status of all LIFs to “up” so the LIFs can accept and send traffic on the node by entering the following command:

```
network interface modify -home-port port_name -home-node node3 -lif
data -status-admin up
```

- d. Enter the following command on either node and examine its output to verify that LIFs have been moved to the correct ports and that the LIFs have the status of “up”:

```
network interface show -home-node node3 -role data
```

- e. If any LIFs are down, set the administrative status of the LIFs to "up" by entering the following command, once for each LIF:

```
network interface modify -vserver vserver_name -lif lif_name -status-
admin up
```

9. Resume the operation to prompt the system to perform the required post-checks:

```
system controller replace resume
```

The system performs the following post-checks:

- Cluster quorum check
- Cluster health check
- Aggregates reconstruction check
- Aggregate status check
- Disk status check
- Cluster LIF status check

## Stage 4 Relocating and retiring node2

---

During Stage 4, all non-root aggregates are relocated from node2 to node3; the non-SAN data LIFs owned by node2 are moved to node3. You must record the necessary node2 information and then retire node2.

### Steps

1. [Relocating non-root aggregates and NAS data LIFs from node2 to node3](#) on page 43
2. [Retiring node2](#) on page 45

## Relocating non-root aggregates and NAS data LIFs from node2 to node3

Before you can replace node2 with node4, you must relocate the non-root aggregates that are owned by node2 to node3.

### Before you begin

After the post-checks are completed from the previous stage, the resource release for node2 starts automatically. The non-root aggregates and non-SAN data LIFs are migrated from node2 to node3.

### About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. Moving SAN LIFs is not necessary for cluster or service health during the upgrade.

**Note:** The home-owner for the aggregates and LIFs are not modified; only the current-owner is modified.

Once the aggregates and LIFs are migrated, the operation is paused for verification purposes. At this stage, you must verify whether or not all the non-root aggregates and non-SAN data LIFs are migrated to node3.

### Steps

1. Verify that all the non-root aggregates are online and their state on node3 by entering the following command:

```
storage aggregate show -node node3 -state online -root false
```

### Example

The following example shows that the non-root aggregates on node2 are online:

```
cluster::> storage aggregate show -node node3 state online -root false
Aggregate Size      Available Used% State #Vols Nodes RAID Status
-----
aggr_1    744.9GB 744.8GB 0%  online 5 node2 raid_dp normal
aggr_2    825.0GB 825.0GB 0%  online 1 node2 raid_dp normal
2 entries were displayed.
```

If the aggregates have gone offline or become foreign on node3, bring them online by using the following command on node3, once for each aggregate:

```
storage aggregate online -aggregate aggr_name
```

2. Verify that all the volumes are online on node3 by entering the following command on node3 and examining its output:

```
volume show -node node3 -state offline
```

If any volumes are offline on node3, bring them online by using the following command on node3, once for each volume:

```
volume online -vserver vservice-name -volume volume-name
```

The *vservice-name* to use with this command is found in the output of the previous `volume show` command.

3. Verify that the LIFs have been moved to the correct ports and have a status of “up”. If any LIFs are down, set the administrative status of the LIFs to “up” by entering the following command, once for each LIF:

```
network interface modify -vserver vservice_name -lif LIF_name -home-node nodename -status-admin up
```

4. If the ports currently hosting data LIFs will not exist on the new hardware, remove them from the broadcast domain:

```
network port broadcast-domain remove-ports
```

5. Enter the following command and examine its output to verify that there are no data LIFs remaining on node2:

```
network interface show -curr-node node2 -role data
```

6. If you have interface groups or VLANs configured, complete the following substeps:

- a. Record VLAN and ifgrp information so you can recreate the VLANs and ifgrps on node3 after node3 is booted up.

- b. Remove the VLANs from the interface groups by entering the following command:

```
network port vlan delete -node nodename -port ifgrp -vlan-id VLAN_ID
```

- c. Enter the following command and examine its output to see if there are any interface groups configured on the node:

```
network port ifgrp show -node node2 -ifgrp ifgrp_name -instance
```

### Example

The system displays interface group information for the node as shown in the following example:

```
cluster::> network port ifgrp show -node node2 -ifgrp a0a -instance
Node: node3
Interface Group Name: a0a
Distribution Function: ip
Create Policy: multimode_lacp
MAC Address: 02:a0:98:17:dc:d4
Port Participation: partial
Network Ports: e2c, e2d
Up Ports: e2c
Down Ports: e2d
```

- d. If any interface groups are configured on the node, record the names of those groups and the ports assigned to them, and then delete the ports by entering the following command, once for each port:

```
network port ifgrp remove-port -node nodename -ifgrp ifgrp_name -port  
netport
```

## Retiring node2

To retire node2, you need to shut node2 down properly and remove it from the rack or chassis.

### Step

1. Resume the operation:

```
system controller replace resume
```

### Result

The node halts automatically.

### After you finish

You can decommission node2 after the upgrade is completed. See [Decommissioning the old system](#) on page 69.

## Stage 5 Installing and booting node4

---

During Stage 5, you install and boot node4, map the cluster and node-management ports from node2 to node4, and verify the node4 installation. You also move the NAS data LIFs owned by node2 from node3 to node4, and relocate node2's aggregates from node3 to node4.

### Steps

1. [Installing and booting node4](#) on page 46
2. [Setting the FC or UTA/UTA2 configuration on node4](#) on page 50
3. [Mapping ports from node2 to node4](#) on page 56
4. [Joining the quorum when a node has a different set of network ports](#) on page 60
5. [Verifying the node4 installation](#) on page 61
6. [Moving non-root aggregates and NAS data LIFs owned by node2 from node3 to node4](#) on page 62

## Installing and booting node4

You need to install node4 in the rack, transfer node2 connections to node4, boot node4, and install ONTAP. You must also reassign any node2 spare disks, any disks belonging to the root volume, and any non-root aggregates that were not relocated to node3 earlier.

### About this task

The relocation operation is paused at the beginning of this stage. This process is largely automated; the operation pauses to allow you to check its status. You must manually resume the operation. In addition, you must verify the NAS data LIFs have successfully moved to node4.

You need to netboot node4 if it does not have the same or a later version of ONTAP 9 that is installed on node2. After you install node4, boot it from the ONTAP 9 image stored on the web server. You can then download the correct files to the boot media device for subsequent system boots.

**Important:** If you are upgrading a V-Series system or a system with FlexArray Virtualization Software that is connected to storage arrays, you need to complete Step [1](#) on page 46 through Step [18](#) on page 49, then leave this section and follow instructions in [Configuring FC ports on node4](#) on page 50 and [Checking and configuring UTA/UTA2 ports on node4](#) on page 52 as needed, entering the commands in Maintenance mode. You then need to return to this section and resume the procedure with Step [20](#) on page 49.

However, if you are upgrading a system with storage disks, you need to complete this entire section and then proceed to the [Setting the FC or UTA/UTA2 configuration on node4](#) on page 50 entering commands at the cluster prompt.

### Steps

1. Make sure that node4 has sufficient rack space.
 

If node4 is in a separate chassis from node2, you can put node4 in the same location as node3. If node2 and node4 are in the same chassis, then node4 is already in its appropriate rack location.
2. Install node4 in the rack, following the instructions in the *Installation and Setup Instructions* for the node model.
3. Cable node4, moving the connections from node2 to node4.

Cable the following connections, using the instructions in the *Installation and Setup Instructions* or the *FlexArray Virtualization Installation Requirements and Reference* for the node4 platform, the appropriate disk shelf guide, and the *ONTAP 9 High-Availability Configuration Guide*:

- Console (remote-management port)
- Cluster ports
- Data ports
- Cluster and node-management ports
- Storage
- SAN configurations: iSCSI Ethernet and FC switch ports

**Note:** You do not need to move the interconnect card/FC\_VI card or interconnect/FC\_VI cable connection from node2 to node4 because most platform models have unique interconnect card models.

4. Turn on the power to node4, and then interrupt the boot by pressing Ctrl-C to access the boot environment prompt.

**Note:** When you boot node4, you might see the following message:

```
WARNING: The battery is unfit to retain data during a power
         outage. This is likely because the battery is
         discharged but could be due to other temporary
         conditions.
         When the battery is ready, the boot process will
         complete and services will be engaged.
         To override this delay, press 'c' followed by 'Enter'
```

5. If you see the warning message above, take the following actions:
  - a. Check for any console messages that might indicate a problem other than a low NVRAM battery and, if necessary, take any required corrective action.
  - b. Allow the battery to charge and the boot process to finish.

**Attention:** Do not override the delay. Failure to allow the battery to charge could result in a loss of data.

**Note:** Refer to [Preparing for netboot](#) on page 22.

6. Configure the netboot connection by choosing one of the following actions.

**Note:** You should use the management port and IP as the netboot connection. Do not use a data LIF IP or a data outage might occur while the upgrade is being performed.

If Dynamic Host Configuration Protocol (DHCP) is...	Then...
Running	Configure the connection automatically by entering the following command at the boot environment prompt:  <code>ifconfig e0M -auto</code>

If Dynamic Host Configuration Protocol (DHCP) is...	Then...
Not running	<p>Manually configure the connection by entering the following command at the boot environment prompt:</p> <pre data-bbox="667 373 1317 432"><b>ifconfig e0M -addr=filer_addr -mask=netmask -gw=gateway -dns=dns_addr domain=dns_domain</b></pre> <p><i>filer_addr</i> is the IP address of the storage system.</p> <p><i>netmask</i> is the network mask of the storage system.</p> <p><i>gateway</i> is the gateway for the storage system.</p> <p><i>dns_addr</i> is the IP address of a name server on your network.</p> <p><i>dns_domain</i> is the Domain Name Service (DNS) domain name. If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL; you need only the server's host name.</p> <p><b>Note:</b> Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p>

7. Perform netboot by entering the following command on node4:

```
netboot http://path_to_the_web-accessible_directory/netboot/kernel
```

The *path\_to\_the\_web-accessible\_directory* should lead to where you downloaded the `netboot.tgz` in Step 1a in [Preparing for netboot](#) on page 22.

**Note:** Do not interrupt the boot.

8. From the boot menu, select option (7) Install new software first.

This menu option downloads and installs the new ONTAP image to the boot device.

**Note:** Disregard the following message: "This procedure is not supported for Non-Disruptive Upgrade on an HA pair". The note applies to nondisruptive upgrades of ONTAP, and not upgrades of controllers.

9. If you are prompted to continue the procedure, enter **y**, and when prompted for the package, enter the URL `http://path_to_the_web-accessible_directory/image.tgz`

10. Complete the following substeps:

- a. Enter **n** to skip the backup recovery when you see the following prompt:

```
Do you want to restore the backup configuration now? {y|n}
n
```

- b. Reboot by entering **y** when you see the following prompt:

```
The node must be rebooted to start using the newly installed
software. Do you want to reboot now? {y|n}
```

The controller module reboots but stops at the boot menu because the boot device was reformatted and the configuration data needs to be restored.

11. Select reboot option 8 from the boot menu and interrupt the autoboot by pressing Ctrl-C.

12. On node3, check the system date, time, and time zone:

```
date
```

13. On node4, check the date at the boot environment prompt:

```
show date
```

14. If necessary, set the date on node4:

```
set date mm/dd/yyyy
```

15. On node4, check the time at the boot environment prompt:

```
show time
```

16. If necessary, set the time on node4:

```
set time hh:mm:ss
```

17. If necessary, set the partner system ID on node4:

```
setenv partner-sysid node3_sysid
```

- a. Save the settings:

```
saveenv
```

18. On the new node, in boot loader, the `partner-sysid` parameter needs to be set. For node4, `partner-sysid` needs to be that of node3. Verify the `partner-sysid` for node3 using the following command:

```
printenv partner-sysid
```

19. Take one of the following actions:

If your system...	Then
Has disks and no back-end storage	Go to Step <a href="#">20</a> on page 49
Is a V-Series system or a system with FlexArray Virtualization software connected to storage arrays	<p><b>a.</b> Go to <a href="#">Setting the FC or UTA/UTA2 configuration on node4</a> on page 50 and complete the subsections in it.</p> <p><b>b.</b> Return to this section and complete the remaining steps, beginning with Step <a href="#">20</a> on page 49.</p> <p><b>Important:</b> You must reconfigure FC onboard ports, CNA onboard ports, and CNA cards before you boot ONTAP on the V-Series or system with FlexArray Virtualization software.</p>

20. Add the FC initiator ports of the new node to the switch zones.

If required, modify the onboard ports to initiator by referring to the [Configuring FC ports on node4](#) on page 50 . See your storage array and zoning documentation for instructions.

21. Add the FC initiator ports to the storage array as new hosts, mapping the array LUNs to the new hosts.

See your storage array and zoning documentation for instructions.

22. Modify the World Wide Port Name (WWPN) values in the host or volume groups associated with array LUNs on the storage array.

Installing a new controller module changes the WWPN values associated with each onboard FC port.

23. If your configuration uses switch-based zoning, adjust the zoning to reflect the new WWPN values.

24. If NetApp Storage Encryption (NSE) is in use on this configuration, the `setenv bootarg.storageencryption.support` command must be set to **true**, and the

`kmip.init.maxwait` variable needs to be set to **off** to avoid a boot loop after the node2 configuration is loaded:

```
setenv bootarg.storageencryption.support true
setenv kmip.init.maxwait off
```

25. Boot node into boot menu by running `boot_ontap menu`.  
If you do not have FC or UTA/UTA2 configuration, execute Step [15](#) on page 55 so that node4 can recognize node2's disks.
26. For V-Series systems and systems with FlexArray Virtualization software connected to storage arrays proceed to Step [15](#) on page 55.

## Setting the FC or UTA/UTA2 configuration on node4

If node4 has onboard FC ports, onboard unified target adapter (UTA/UTA2) ports, or a UTA/UTA2 card, you must configure the settings before completing the rest of the procedure.

### About this task

You might need to complete the [Configuring FC ports on node4](#) on page 50 section, the [Checking and configuring UTA/UTA2 ports on node4](#) on page 52 section, or both sections.

If node4 does not have onboard FC ports, onboard UTA/UTA2 ports, or a UTA/UTA2 card, and you are upgrading a system with storage disks, you can skip to the [Mapping ports from node2 to node4](#) on page 56 section. However, if you have a V-Series system or have FlexArray Virtualization Software and are connected to storage arrays, and node4 does not have onboard FC ports, onboard UTA/UTA2 ports, or a UTA/UTA2 card, you must return to the [Installing and booting node4](#) on page 46 section and resume at Step [19](#) on page 49. Make sure that node4 has sufficient rack space. If node4 is in a separate chassis from node2, you can put node4 in the same location as node3. If node2 and node4 are in the same chassis, then node4 is already in its appropriate rack location.

### Choices

- [Configuring FC ports on node4](#) on page 50
- [Checking and configuring UTA/UTA2 ports on node4](#) on page 52

## Configuring FC ports on node4

If node4 has FC ports, either onboard or on an FC adapter, you must set port configurations on the node before you bring it into service because the ports are not preconfigured. If the ports are not configured, you might experience a disruption in service.

### Before you begin

You must have the values of the FC port settings from node2 that you saved in the section [Preparing the nodes for upgrade](#) on page 13.

### About this task

You can skip this section if your system does not have FC configurations. If your system has onboard UTA/UTA2 ports or a UTA/UTA2 adapter, you configure them in [Checking and configuring UTA/UTA2 ports on node4](#) on page 52.

**Important:** If your system has storage disks, you must enter the commands in this section at the cluster prompt. If you have a V-Series system or a system with FlexArray Virtualization Software connected to storage arrays, you enter commands in this section in Maintenance mode.

## Steps

1. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	Enter the following command:  <b>system node hardware unified-connect show</b>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Enter the following command:  <b>ucadmin show</b>

The system displays information about all FC and converged network adapters on the system.

2. Compare the FC settings on node4 with the settings that you captured earlier from node1.
3. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	Modify the FC ports on node4 as needed by entering one of the following commands: <ul style="list-style-type: none"> <li>• To program target ports: <b>system node hardware unified-connect modify -type   -t target -adapter port_name</b></li> <li>• To program initiator ports: <b>system node unified-connect modify -type   -t initiator -adapter port_name</b></li> </ul> -type is the FC4 type, target or initiator.
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Modify the FC ports on node4 as needed by entering the following command: <b>ucadmin modify -m fc -t initiator -f adapter_port_name</b> -t is the FC4 type, target or initiator.  <b>Note:</b> The FC ports need to be programmed as initiators.

4. Exit Maintenance mode by entering the following command:  
**halt**
5. After you enter the command, wait until the system stops at the boot environment prompt.
6. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	<ul style="list-style-type: none"> <li>• Skip this section and go to <a href="#">Mapping ports from node2 to node4</a> on page 56 if node4 does not have a UTA/UTA2 card or UTA/UTA2 onboard ports.</li> </ul>

If the system that you are upgrading...	Then...
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<ul style="list-style-type: none"> <li>Go to Checking and configuring UTA/UTA2 ports on node4 if node4 has a UTA/UTA2 card or UTA/UTA2 onboard ports.</li> <li>Skip the section <a href="#">Checking and configuring UTA/UTA2 ports on node4</a> on page 52 if node4 does not have a UTA/UTA2 card or UTA/UTA2 onboard ports, return to the section <a href="#">Installing and booting node4</a> on page 46, and resume the section at Step <a href="#">20</a> on page 49.</li> </ul>

## Checking and configuring UTA/UTA2 ports on node4

If node4 has onboard UTA/UTA2 ports or a UTA/UTA2A card, you must check the configuration of the ports and configure them, depending on how you want to use the upgraded system.

### Before you begin

You must have the correct SFP+ modules for the UTA/UTA2 ports.

### About this task

UTA/UTA2 ports can be configured into native FC mode or UTA/UTA2A mode. FC mode supports FC initiator and FC target; UTA/UTA2 mode allows concurrent NIC and FCoE traffic to share the same 10GbE SFP+ interface and supports FC target.

**Note:** NetApp marketing materials might use the term UTA2 to refer to CNA adapters and ports. However, the CLI uses the term CNA.

UTA/UTA2 ports might be on an adapter or on the controller with the following configurations:

- UTA/UTA2 cards ordered at the same time as the controller are configured before shipment to have the personality you requested.
- UTA/UTA2 cards ordered separately from the controller are shipped with the default FC target personality.
- Onboard UTA/UTA2 ports on new controllers are configured (before shipment) to have the personality you requested.

However, you should check the configuration of the UTA/UTA2 ports on node4 and change it, if necessary.

**Important:** If your system has storage disks, you enter the commands in this section at the cluster prompt unless directed to enter Maintenance mode. If you have a V-Series system or a system with FlexArray Virtualization software that is connected to storage arrays, you enter commands in this section in Maintenance mode. You must be in Maintenance mode to configure UTA/UTA2 ports.

### Steps

1. Check how the ports are currently configured by entering one of the following commands on node4:

If the system...	Then...
Has storage disks	Enter the following command:  <b>system node hardware unified-connect show</b>

If the system...	Then...
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Enter the following command: <b>ucadmin show</b>

**Example**

The system displays output similar to the following examples:

```
cluster1::> system node hardware unified-connect show
Node   Adapter  Current Mode   Current Type   Pending Mode   Pending Type   Admin Status
-----
f-a    0e       fc     initiator -      -      online
f-a    0f       fc     initiator -      -      online
f-a    0g       cna    target  -      -      online
f-a    0h       cna    target  -      -      online
f-b    0e       fc     initiator -      -      online
f-b    0f       fc     initiator -      -      online
f-b    0g       cna    target  -      -      online
f-b    0h       cna    target  -      -      online
8 entries were displayed.
```

```
*> ucadmin show
Node   Adapter  Current Mode   Current Type   Pending Mode   Pending Type   Admin Status
-----
f-a    0e       fc     initiator -      -      online
f-a    0f       fc     initiator -      -      online
f-a    0g       cna    target  -      -      online
f-a    0h       cna    target  -      -      online
f-a    0e       fc     initiator -      -      online
f-a    0f       fc     initiator -      -      online
f-a    0g       cna    target  -      -      online
f-a    0h       cna    target  -      -      online
*>
```

2. If the current SFP+ module does not match the desired use, replace it with the correct SFP+ module. Contact your NetApp representative to obtain the correct SFP+ module.
3. Examine the output of the `system node hardware unified-connect show` or `ucadmin show` command and determine whether the UTA/UTA2 ports have the personality you want.

4. Take one of the following actions:

If the CNA ports...	Then...
Do not have the personality that you want	Go to Step 5.
Have the personality that you want	Skip Step 5 through Step 12 and go Step 13.

5. Take one of the following actions:

If you are configuring...	Then...
Ports on a UTA/UTA2A card	Go to Step 7.
Onboard UTA/UTA2 ports	Skip Step 7 and go to Step 8.

6. If the adapter is in initiator mode, and if the UTA/UTA2 port is online, take the UTA/UTA2 port offline by entering the following command:

```
storage disable adapter adapter_name
```

Adapters in target mode are automatically offline in Maintenance mode.

- If the current configuration does not match the desired use, enter the following command to change the configuration as needed:

```
ucadmin modify -m fc|cna -t initiator|target adapter_name
```

- m is the personality mode: FC or 10GbE UTA.
- t is the FC4 type: target or initiator.

**Note:** You need to use FC initiator for tape drives and FlexArray Virtualization systems. You need to use the FC target for SAN clients.

- Verify the settings by entering the following command and examining its output:

```
ucadmin show
```

- Verify the settings by entering one of the following commands:

If the system...	Then...
Has storage disks	Enter the following command: <b>system node hardware unified-connect show</b>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Enter the following command: <b>ucadmin show</b>

### Example

The output in the following examples shows that the FC4 type of adapter 1b is changing to initiator and that the mode of adapters 2a and 2b is changing to cna.

```
cluster1::> system node hardware unified-connect show
Node  Adapter  Current Mode  Current Type  Pending Mode  Pending Type  Admin Status
-----
f-a   1a        fc      initiator -      -      initiator online
f-a   1b        fc      target  -      initiator online
f-a   2a        fc      target  cna    -      online
f-a   2b        fc      target  cna    -      online
4 entries were displayed.
```

```
*> ucadmin show
Node  Adapter  Current Mode  Current Type  Pending Mode  Pending Type  Admin Status
-----
f-a   1a        fc      initiator -      -      initiator online
f-a   1b        fc      target  -      initiator online
f-a   2a        fc      target  cna    -      online
f-a   2b        fc      target  cna    -      online
4 entries were displayed.
*>
```

- Place any target ports online by entering one of the following commands, once for each port:

If the system...	Then...
Has storage disks	<b>network fcp adapter modify -node node_name -adapter adapter_name -state up</b>

If the system...	Then...
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<b>fcpx config adapter_name up</b>

11. Cable the port.

12. Take one of the following actions:

If the system...	Then...
Has storage disks	Go to the section <a href="#">Mapping ports from node2 to node4</a> on page 56.
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Return to the section <a href="#">Installing and booting node4</a> on page 46 and resume the section at Step 20 on page 49.

13. Exit Maintenance mode by entering the following command:

```
halt
```

14. Boot node into boot menu by running

```
boot_ontap menu
```

```
.
```

15. On node4, go to the boot menu, and using 22/7, select the hidden option `boot_after_controller_replacement`. At the prompt, enter `node2` to reassign the disks of node2 to node4, as per the following example.

### Example

```

LOADER-A> boot_ontap menu
...
*****
*                                     *
* Press Ctrl-C for Boot Menu. *
*                                     *
*****
.
Please choose one of the following:

(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
Selection (1-9)? 22/7
.
.
(boot_after_controller_replacement) Boot after controller upgrade
(9a) Unpartition all disks and remove their ownership information.
(9b) Clean configuration and initialize node with partitioned disks.
(9c) Clean configuration and initialize node with whole disks.
(9d) Reboot the node.
(9e) Return to main boot menu.

Please choose one of the following:

(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
Selection (1-9)? boot_after_controller_replacement
.
This will replace all flash-based configuration with the last backup to
disks. Are you sure you want to continue?: yes
.
.
Controller Replacement: Provide name of the node you would like to replace: <name of the node being
replaced>
.

```

```

.
Changing sysid of node <node being replaced> disks.
Fetched sanown old_owner_sysid = 536953334 and calculated old sys id = 536953334
Partner sysid = 4294967295, owner sysid = 536953334
.
.
.
Terminated
<node reboots>
.
.
System rebooting...
.
Restoring env file from boot media...
copy_env_file:scenario = head upgrade
Successfully restored env file from boot media...
.
.
System rebooting...
.
.
WARNING: System ID mismatch. This usually occurs when replacing a boot device or NVRAM cards!
Override system ID? {y|n} y
Login:
...

```

16. If the system goes into a reboot loop with the message `no disks found`, this is because it has reset the ports back to the target mode and therefore is unable to see any disks. Continue with Steps 17 on page 56 to 22 on page 56 to resolve this.
17. Press Ctrl-C and break the loop.
18. In loader prompt, use the command `boot_ontap menu` and enter maintenance mode.
19. In maintenance mode, use the command `ucadmin show` to display all the previously set initiator ports that are now in target mode. Change the ports back to initiator mode using the following command
 

```
ucadmin modify -m fc -t initiator -f <adapter name>
```
20. Enter `ucadmin show` to verify that the ports have been changed to initiator mode.
21. Exit maintenance mode by entering the command `halt`.
22. In loader prompt, use the command `boot_ontap`.

Now, on booting, the node is able to detect all the disks that were previously assigned to it and can boot up as expected.

## Mapping ports from node2 to node4

You must verify that the physical ports on node2 map correctly to the physical ports on node4, which will allow node4 to communicate with other nodes in the cluster and with the network after the upgrade.

### About this task

Capture information about the ports on the new nodes from the *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com). You will use the information later in this section.

The software configuration of node4 must match the physical connectivity of node4 and IP connectivity must be restored before you continue with the upgrade.

Port settings might vary, depending on the model of the nodes. You must make the original node's port and LIF configuration compatible with what you plan the new node's configuration to be. This is because the new node replays the same configuration when it boots, meaning when you boot node4 that Data ONTAP will try to host LIFs on the same ports that were used on node2.

Therefore, if the physical ports on node2 do not map directly to the physical ports on node4, then software configuration changes will be required to restore cluster, management, and network

connectivity after the boot. In addition, if the cluster ports on node2 do not directly map to the cluster ports on node4, node4 might not automatically rejoin the quorum when it is rebooted until a software configuration change is made to host the cluster LIFs on the correct physical ports.

**Steps**

1. Record all the node2 cabling information for node2, the ports, broadcast domains, and IPspaces, in the table:

LIF	Node 2 ports	Node2 IPspaces	Node2 broadcast domains	Node4 ports	Node4 IPspaces	Node4 broadcast domains
Cluster 1						
Cluster 2						
Cluster 3						
Cluster 4						
Node management						
Cluster management						
Data 1						
Data 2						
Data 3						
Data 4						
SAN						
Intercluster port						

2. Record all the cabling information for node4, the ports, broadcast domains, and IPspaces, in the table.
3. Follow these steps to verify if the setup is a two-node switchless cluster:
  - a. Set the privilege level to advanced:

```
cluster::> set -privilege advanced
```

- b. Verify if the setup is a two-node switchless cluster:

```
cluster::> network options switchless-cluster show
```

**Example**

```
cluster::*> network options switchless-cluster show
Enable Switchless Cluster: false/true
```

The value of this command must match the physical state of the system.

- c. Return to the administration privilege level:

```
cluster::*> set -privilege admin
cluster::>
```

4. Follow these steps to place node4 into quorum:

- a. Boot node4. See [Installing and booting node4](#) on page 46 to boot the node if you have not already done so.
- b. Verify that the new cluster ports are in the Cluster broadcast domain:

```
network port show -node node -port port -fields broadcast-domain
```

### Example

The following example shows that port e0a is in the Cluster domain on node4:

```
cluster::> network port show -node node4 -port e0a -fields
broadcast-domain

node      port broadcast-domain
-----  ---  -----
node4     e1a  Cluster
```

- c. If the cluster ports are not in the Cluster broadcast-domain, add them with the following command:

```
broadcast-domain add-ports -ip-space Cluster -broadcast-domain Cluster
-ports node:port
```

- d. Add the correct ports to the Cluster broadcast domain:

```
network port modify -node -port -ip-space Cluster -mtu 9000
```

### Example

This example adds Cluster port e1b on node4:

```
network port modify -node node4 -port e1b -ip-space Cluster -mtu 9000
```

**Note:** For a MetroCluster configuration, you may not be able to change the broadcast domain of a port because it is associated with a port hosting the LIF of a sync-destination SVM and see errors similar to, but not restricted, to the following:

```
command failed: This operation is not permitted on a Vserver that
is
    configured as the destination of a MetroCluster Vserver
relationship.
```

Enter the following command from the corresponding sync-source SVM on the remote site to reallocate the sync-destination LIF to an appropriate port:

```
metrocluster vserver resync -vserver vserver_name
```

- e. Migrate the cluster LIFs to the new ports, once for each LIF:

```
network interface migrate -vserver Cluster -lif lif_name -source-node
node4 -destination-node node4 -destination-port port_name
```

- f. Modify the home port of the cluster LIFs:

```
network interface modify -vserver Cluster -lif lif_name -home-port
port_name
```

- g. Remove the old ports from the Cluster broadcast domain:

```
network port broadcast-domain remove-ports
```

### Example

This command removes port e0d on node4:

```
network port broadcast-domain remove-ports -ip-space Cluster -
broadcast-domain Cluster -ports node4:e0d
```

- h. Verify that node4 has rejoined quorum:

```
cluster show -node node4 -fields health
```

5. Adjust the broadcast domains hosting your cluster LIFs and node-management/cluster-management LIFs. Ensure that each broadcast domain contains the correct ports. A port cannot be moved between broadcast domains if it is hosting or is home to a LIF so you may need to migrate and modify the LIFs as shown in the following steps:

- a. Display the home port of a LIF:

```
network interface show -fields home-node,home-port
```

- b. Display the broadcast domain containing this port:

```
network port broadcast-domain show -ports node_name:port_name
```

- c. Add or remove ports from broadcast domains:

```
network port broadcast-domain add-ports
```

```
network port broadcast-domain remove-ports
```

- d. Modify a LIF's home port:

```
network interface modify -vserver vservers -lif lif_name -home-port
port_name
```

6. Adjust the intercluster broadcast domains and migrate the intercluster LIFs, if necessary, using the same commands shown in Step 5.
7. Adjust any other broadcast domains and migrate the data LIFs, if necessary, using the same commands shown in Step 5.
8. If there were any ports on node2 that no longer exist on node4, follow these steps to delete them:

- a. Access the advanced privilege level on either node:

```
set -privilege advanced
```

- b. To delete the ports:

```
network port delete -node node_name -port port_name
```

- c. Return to the admin level:

```
set -privilege admin
```

9. Adjust all the LIF failover groups:

```
network interface modify -failover-group failover_group -failover-policy
failover_policy
```

### Example

The following command sets the failover policy to broadcast-domain-wide and uses the ports in failover group fg1 as failover targets for LIF data1 on node4:

```
network interface modify -vserver node4 -lif data1 failover-policy
broadcast-domain-wide -failover-group fg1
```

See “Configuring failover settings on a LIF” in the *ONTAP 9 Network Management Guide* or *ONTAP 9 Commands: Manual Page Reference* for more information.

10. Verify the changes on node4:

```
network port show -node node4
```

## Joining the quorum when a node has a different set of network ports

The node with the new controller boots and attempts to join the cluster automatically at first; however if the new node has a different set of network ports, you must perform the following steps to ensure that the node successfully joins the quorum.

### About this task

You can use these instructions for any relevant node. Node3 is used throughout the following sample.

### Steps

1. Verify that the new cluster ports are in the Cluster broadcast domain by entering the following command and checking its output:

```
network port show -node node -port port -fields broadcast-domain
```

### Example

The following example shows that port e1a is in the Cluster domain on node3:

```
cluster::> network port show -node node3 -port e1a -fields broadcast-
domain
node   port   broadcast-domain
-----
node3  e1a   Cluster
```

2. Add the correct ports to the Cluster broadcast domain by entering the following command and checking its output:

```
network port modify -node -port -ipspace Cluster -mtu 9000
```

### Example

This example adds Cluster port e1b on node3:

```
network port modify -node node3 -port e1b -ipspace Cluster -mtu 9000
```

3. Migrate the cluster LIFs to the new ports, once for each LIF, using the following command:

```
network interface migrate -vserver Cluster -lif lif_name -source-node
node3 -destination-node node3 -destination-port port_name
```

4. Modify the home port of the cluster LIFs as follows:

```
network interface modify -vserver Cluster -lif lif_name -home-port
port_name
```

5. If the cluster ports are not in the Cluster broadcast-domain, add them with the following command:

```
network port broadcast-domain add-ports -ipspace Cluster -
broadcastdomain Cluster -ports node:port
```

6. Remove the old ports from the Cluster broadcast domain. You can use for any relevant node. The following command removes port e0d on node3:
 

```
network port broadcast-domain remove-ports network port broadcast-domain
remove-ports -ip-space Cluster -broadcast-domain Cluster -ports node3:e0d
```
7. Verify the node has rejoined quorum as follows:
 

```
cluster show -node node3 -fields health
```
8. Adjust the broadcast domains hosting your cluster LIFs and node-management/cluster management LIFs. Ensure that each broadcast domain contains the correct ports. A port cannot be moved between broadcast domains if it is hosting or is home to a LIF, so you might need to migrate and modify the LIFs as follows:
  - a. Display the home port of a LIF:
 

```
network interface show -fields home-node,home-port
```
  - b. Display the broadcast domain containing this port:
 

```
network port broadcast-domain show -ports node_name:port_name
```
  - c. Add or remove ports from broadcast domains:
 

```
network port broadcast-domain add-ports network port broadcast-domain
remove-port
```
  - d. Modify a home port of a LIF:
 

```
network interface modify -vserver vservers -lif lif_name -home-port
port_name
```

 Adjust the intercluster broadcast domains and migrate the intercluster LIFs, if necessary. The data LIFs remain unchanged.

## Verifying the node4 installation

After you install and boot node4, you need to verify that it is installed correctly, that it is part of the cluster, and that it can communicate with node3.

### About this task

At this point in the procedure, the operation will have paused as node4 joins quorum.

### Steps

1. Verify that node4 has joined quorum:
 

```
cluster show -node node4 -fields health
```
2. Verify that node4 is part of the same cluster as node3 and healthy by entering the following command:
 

```
cluster show
```
3. Check the status of the operation and verify that the configuration information for node4 is the same as node2:
 

```
system controller replace show-details
```

If the configuration is different for node4, a system disruption might occur later in the procedure.

## Re-creating VLANs, ifgrps, and broadcast domains on node4

After you confirm that node4 is in quorum and can communicate with node3, you must re-create node2's VLANs, ifgrps, and broadcast domains on node4. You must also add the node3 ports to the newly re-created broadcast domains.

### About this task

For more information on creating and re-creating VLANs, ifgrps, and broadcast domains, see the *ONTAP 9.5 Cluster-Mode Administration Reference*.

### Steps

1. Re-create the VLANs on node4 using the node2 information recorded in the [Relocating non-root aggregates and NAS data LIFs from node2 to node3](#) on page 43 section:

```
network port vlan create -node node4 -vlan vlan-names
```

2. Re-create the ifgrps on node4 using the node2 information recorded in the [Relocating non-root aggregates and NAS data LIFs from node2 to node3](#) on page 43 section:

```
network port ifgrp create -node node4 -ifgrp port_ifgrp_names-distr-func
```

3. Re-create the broadcast domains on node4 using the node2 information recorded in the [Relocating non-root aggregates and NAS data LIFs from node2 to node3](#) on page 43 section:

```
network port broadcast-domain create -ip-space Default -broadcast-domain broadcast_domain_names -mtu mtu_size -ports node_name:port_name,node_name:port_name
```

4. Add the node4 ports to the newly re-created broadcast domains:

```
network port broadcast-domain add-ports -broadcast-domain broadcast_domain_names -ports node_name:port_name,node_name:port_name
```

## Moving non-root aggregates and NAS data LIFs owned by node2 from node3 to node4

After you verify the node4 installation and before you relocate aggregates from node3 to node4, you need to move the NAS data LIFs belonging to node2 that are currently on node3 from node3 to node4. You also need to verify the SAN LIFs exist on node4.

### About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. Moving SAN LIFs is not necessary for cluster or service health during the upgrade. SAN LIFs are not moved unless they need to be mapped to new ports. You will verify that the LIFs are healthy and located on appropriate ports after you bring node4 online.

### Steps

1. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following tasks:

- Cluster quorum check
- System ID check

- Image version check
- Target platform check
- Network reachability check

The operation pauses at this stage in the network reachability check.

2. Manually verify that the network and all VLANs, ifgrps, and broadcast domains have been configured correctly.
3. Resume the relocation operation:

```
system controller replace resume
```

```
To complete the "Network Reachability" phase, ONTAP network configuration must be manually adjusted to match the new physical network configuration of the hardware. This includes assigning network ports to the correct broadcast domains, creating any required ifgrps and VLANs, and modifying the home-port parameter of network interfaces to the appropriate ports. Refer to the "Using aggregate relocation to upgrade controller hardware on a pair of nodes running ONTAP 9.x" documentation, Stages 3 and 5. Have all of these steps been manually completed? [y/n]
```

4. Enter *y* to continue.
5. The system performs the following checks:

- Cluster health check
- Disk status check
- Cluster LIF status check

After performing the three checks, the system relocates the non-root aggregates and NAS data LIFs owned by node2 to the new controller, node4.

The system pauses once the resource relocation is complete.

6. Check the status of the aggregate relocation and NAS data LIF move operations:

```
system controller replace show-details
```

7. Manually verify that the non-root aggregates and NAS data LIFs have been successfully relocated to node4.

If any aggregates fail to relocate or are vetoed, you must manually relocate the aggregates, or override either the vetoes or destination checks, if necessary. See the section [Relocating failed or vetoed aggregates](#) on page 21 for more information.

8. Ensure that the SAN LIFs are on the correct ports on node4 by completing the following substeps:

- a. Enter the following command and examine its output:

```
network interface show -data-protocol iscsi|fc -home-node node4
```

### Example

The system returns output similar to the following example:

```
cluster::> net int show -data-protocol iscsi|fc -home-node node3
      Logical   Status   Network           Current Current  Is
Vserver Interface Admin/Oper Address/Mask      Node   Port   Home
-----
vs0
  a0a         up/down  10.63.0.53/24    node3  a0a   true
  data1       up/up    10.63.0.50/18    node3  e0c   true
  rads1       up/up    10.63.0.51/18    node3  e1a   true
```

vs1	rads2	up/down	10.63.0.52/24	node3	e1b	true
	lif1	up/up	172.17.176.120/24	node3	e0c	true
	lif2	up/up	172.17.176.121/24	node3	e1a	true

- b. If node4 has any SAN LIFs or groups of SAN LIFs that are on a port that did not exist on node2 or that need to be mapped to a different port, move them to an appropriate port on node4 by completing the following substeps:

- i. Set the LIF status to down by entering the following command:

```
network interface modify -vserver vs1 -lif lif_name -status-admin down
```

- ii. Remove the LIF from the port set:

```
portset remove -vserver vs1 -portset portset_name -port-name port_name
```

- iii. Enter one of the following commands:

- Move a single LIF by entering the following command:

```
network interface modify -vserver vs1 -lif lif_name -home-port new_home_port
```

- Move all the LIFs on a single nonexistent or incorrect port to a new port by entering the following command:

```
network interface modify {-home-port port_on_node1 -home-node node1 -role data} -home-port new_home_port_on_node3
```

- Add the LIFs back to the port set:

```
portset add -vserver vs1 -portset portset_name -port-name port_name
```

**Note:** You need to ensure that you move SAN LIFs to a port that has the same link speed as the original port.

- c. Modify the status of all LIFs to “up” so the LIFs can accept and send traffic on the node by entering the following command:

```
network interface modify -home-port port_name -home-node node4 -lif data -status-admin up
```

- d. Enter the following command and examine its output to verify that LIFs have been moved to the correct ports and that the LIFs have the status of “up” by entering the following command on either node and examining the output:

```
network interface show -home-node node4 -role data
```

- e. If any LIFs are down, set the administrative status of the LIFs to up by entering the following command, once for each LIF:

```
network interface modify -vserver vs1 -lif lif_name -status-admin up
```

9. Resume the operation to prompt the system to perform the required post-checks:

```
system controller replace resume
```

The system performs the following post-checks:

- Cluster quorum check
- Cluster health check
- Aggregates reconstruction check

- Aggregate status check
- Disk status check
- Cluster LIF status check

## Stage 6 Completing the upgrade

---

During stage 6, ensure that the new nodes are set up correctly. If one of the new nodes has a unified target adapter, you must restore any port configurations and might need to change the personality of the adapter. You also should set up Storage Encryption if the new nodes are encryption-enabled. You also should decommission the old nodes.

### Steps

1. [Ensuring that the new controllers are set up correctly](#) on page 66
2. [Setting up Storage Encryption on the new controller module](#) on page 68
3. [Decommissioning the old system](#) on page 69
4. [Resuming SnapMirror operations](#) on page 69

## Ensuring that the new controllers are set up correctly

To ensure correct setup, you must enable the HA pair. You must also verify that node3 and node4 can access each other's storage and that neither owns data LIFs belonging to other nodes on the cluster. In addition, you must ensure that node3 owns node1's aggregates and that node4 owns node2's aggregates, and that the volumes for both nodes are online.

### Steps

1. After the post-checks of node2, the storage failover and cluster HA, for the node2 cluster, are enabled. When the operation is done both of the nodes list as completed and the system performs some clean up operations.
2. Verify that storage failover is enabled by entering the following command:

```
storage failover show
```

### Example

The following example shows the output of the command when storage failover is enabled:

```
cluster::> storage failover show
Node           Partner           Takeover
-----         -
node3          node4             true
node4          node3             true
State Description
-----
node3          Connected to node4
node4          Connected to node3
```

3. Verify that node3 and node4 belong to the same cluster by entering the following command and examining the output:
 

```
cluster show
```
4. Verify that node3 and node4 can access each other's storage by entering the following command and examining the output:
 

```
storage failover show -fields local-missing-disks,partner-missing-disks
```
5. Verify that neither node3 nor node4 owns data LIFs home-owned by other nodes in the cluster by entering the following command and examining the output:
 

```
network interface show
```

If either node3 or node4 owns data LIFs home-owned by other nodes in the cluster, use the `network interface revert` command to revert the data LIFs to their home-owner.

- Verify that node3 owns the aggregates from node1 and that node4 owns the aggregates from node2 by entering the following commands:

```
storage aggregate show -owner-name node3
storage aggregate show -owner-name node4
```

- Determine whether any volumes are offline by entering the following commands:

```
volume show -node node3 -state offline
volume show -node node4 -state offline
```

- If any volumes are offline, compare them with the list of offline volumes that you captured in the section [Preparing the nodes for upgrade](#) on page 13, and bring any online any of the offline volumes, as required, by entering the following command, once for each volume:

```
volume online -vserver vsserver-name -volume volume_name
```

- Install new licenses for the new nodes by entering the following commands for each node:

```
system license add -license-code
license_code,license_code,license_code...
```

The license-code parameter accepts a list of 28 upper-case alphabetic character keys. You can add one license at a time, or you can add or multiple licenses at once, each license key separated by a comma.

- To remove all of the old licenses from the original nodes, enter one of the following commands:

```
system license clean-up -unused -expired
system license delete -serial-number node_serial_number -package
licensable_package
```

- To delete all expired licenses, enter:

```
system license clean-up -expired
```

- To delete all unused licenses, enter:

```
system license clean-up -unused
```

- To delete a specific license from a cluster, enter the following commands on the nodes:

```
system license delete -serial-number <node1 serial number> -package *
system license delete -serial-number <node2 serial number> -package *
```

The following output is displayed:

```
Warning: The following licenses will be removed:
<list of each installed package>
Do you want to continue? {y|n}: y
```

Enter `y` to remove all of the packages.

- Verify that the licenses are properly installed by entering the following command and examining its output:

```
system license show
```

You might want to compare the output with the output that you captured in the section [Preparing the nodes for upgrade](#) on page 13 section.

- Configure the SPs by performing the following command on both nodes:

```
system service-processor network modify -node node_name
```

See the *System Administration Reference* for information about the SPs and the *ONTAP 9 Commands: Manual Page Reference* for detailed information about the `system node service-processor network modify` command.

13. Take the following actions on one of the new nodes:
  - a. Enter advanced privilege level by entering the following command:
 

```
set -privilege advanced
```
  - b. Enter the following command:
 

```
storage failover modify -node node_name -cifs-ndo-duration default|medium|low
```

    - Enter **medium** if the system will have workloads in which 50 percent to 75 percent of the operations will be 4 KB or smaller.
    - Enter **low** if the system will have workloads in which 75 percent to 100 percent of the operations will be 4 KB or smaller.
  - c. Return to the admin level by entering the following command:
 

```
set -privilege admin
```
  - d. Reboot the system to ensure that the changes take effect.
14. If you want to set up a switchless cluster on the new nodes, follow the instructions in *Transitioning to a two-node switchless cluster* on the NetApp Support Site.

#### After you finish

If Storage Encryption is enabled on node3 and node4, complete the section [Setting up Storage Encryption on the new controller module](#) on page 68 . Otherwise, complete the section [Decommissioning the old system](#) on page 69.

## Setting up Storage Encryption on the new controller module

If the replaced controller or the HA partner of the new controller uses Storage Encryption, you must configure the new controller module for Storage Encryption, including installing SSL certificates and setting up key management servers.

#### About this task

This procedure includes steps that are performed on the new controller module. You must enter the command on the correct node.

#### Steps

1. Verify that the key management servers are still available, their status, and their authentication key information:
 

```
security key-manager show -status
```

```
security key-manager query
```
2. Add the key management servers listed in the previous step to the key management server list in the new controller:
  - a. Add the key management server:
 

```
security key-manager -add key_management_server_ip_address
```
  - b. Repeat the previous step for each listed key management server.

You can link up to four key management servers.

- c. Verify that the key management servers were added successfully:

```
security key-manager show
```

3. On the new controller module, run the key management setup wizard to set up and install the key management servers.

You must install the same key management servers that are installed on the existing controller module.

- a. Launch the key management server setup wizard on the new node:

```
security key-manager setup -node new_controller_name
```

- b. Complete the steps in the wizard to configure key management servers.

4. Restore authentication keys from all linked key management servers to the new node:

```
security key-manager restore -node new_controller_name
```

## Decommissioning the old system

After upgrading, you can decommission the old system through the NetApp Support Site. Decommissioning the system tells NetApp that the system is no longer in operation and removes it from support databases.

### Steps

1. Navigate to the NetApp Support Site at [mysupport.netapp.com](https://mysupport.netapp.com) and log in.
2. Select **Products > My Products** from the menu.
3. On the **View Installed Systems** page, choose which **Selection Criteria** you want to use to display information about your system.

You can choose one of the following to locate your system:

- Serial Number (located on the back of the unit)
- Serial Numbers for My Location

4. Click **Go!**

A table displays cluster information, including serial numbers.

5. Locate the cluster in the table and select **Decommission this system** from the Product Tool Set drop-down menu.

## Resuming SnapMirror operations

You can resume SnapMirror transfers that were quiesced before upgrade and resume the storage virtual machine (SVM) disaster recovery relationship. The updates are on schedule once the upgrade is completed.

### Steps

1. Verify the storage virtual machine (SVM) disaster recovery SnapMirror state on the destination:

```
snapmirror show
```

2. Resume the SnapMirror relationship:

```
snapmirror resume -destination-vserver
```

## Troubleshooting

---

You might encounter a failure while upgrading the node pair. The node might crash, aggregates might not relocate, or LIFs might not migrate. The cause of the failure and its solution depend on when the failure occurred during the upgrade procedure.

Refer to the table describing the different phases of the procedure in the section [Overview of the ARL upgrade](#) on page 10. The information about failures that can occur is listed by the phase of the procedure.

### Aggregate relocation failures

Aggregate relocation (ARL) might fail at different points during the upgrade.

#### Checking for aggregate relocation failure

During the procedure, ARL might fail in Stage 2, Stage 3, or Stage 5.

##### Steps

1. Enter the following command: `storage aggregate relocation show` and examine its output.

The `storage aggregate relocation show` command shows you which aggregates were successfully relocated and which ones were not, along with the causes of failure.

2. Check the console for any EMS messages.
3. Take one of the following actions:
  - Take the appropriate corrective action, depending on the output of the `storage aggregate relocation show` command and the output of the EMS message.
  - Force relocation of the aggregate or aggregates by using the `override-vetoes` option or the `override-destination-checks` option of the `storage aggregate relocation start` command.

For detailed information about the `storage aggregate relocation start`, `override-vetoes`, and `override-destination-checks` options, see the *ONTAP 9 Commands: Manual Page Reference*.

### Aggregates originally on node1 are owned by node4 after completion of the upgrade

At the end of the upgrade procedure, node3 should be the new home node of aggregates that originally had node1 as the home node. You can relocate them after the upgrade.

##### About this task

Aggregates might fail to relocate properly, having node1 as their home node instead of node3 under the following circumstances:

- During Stage 3, when aggregates are relocated from node2 to node3. Some of the aggregates being relocated have node1 as their home node. For example, such an aggregate could be called `aggr_node_1`. If relocation of `aggr_node_1` fails during Stage 3, and relocation cannot be forced, then the aggregate will be left behind on node2.

- After Stage 4, when node2 is replaced with node4.  
When node2 is replaced, `aggr_node_1` will come online with node4 as its home node instead of node3.

You can fix the incorrect ownership problem after Stage 6 once storage failover has been enabled by completing the following steps:

### Steps

1. Enter the following command to get a list of aggregates:

```
storage aggregate show -nodes node4 -is-home true
```

To identify aggregates that were not correctly relocated, refer to the list of aggregates with the home owner of node1 that you obtained in the section *Preparing the nodes for upgrade* on page 13 and compare it with output of the above command.

2. Compare the output of Step 1 with the output you captured for node1 in the section *Preparing the nodes for upgrade* on page 13 and note any aggregates that were not correctly relocated.

3. Relocate the aggregates left behind on node4 by entering the following command:

```
storage aggregate relocation start -node node4 -aggr aggr_node_1 -
destination node3
```

Do not use the `-ndo-controller-upgrade` parameter during this relocation.

4. Enter the following command to verify that node3 is now the home owner of the aggregates:

```
storage aggregate show -aggregate aggr1,aggr2,aggr3... -fields home-name
```

`aggr1,aggr2,aggr3...` is the list of aggregates that had node1 as the original home owner.

Aggregates that do not have node3 as home owner can be relocated to node3 using the same relocation command in Step 3.

## Reboots, panics, or power cycles

The system might crash – reboot, panic or go through a power cycle – during different stages of the upgrade. The solution to these problems depends on when they occur.

### Reboots, panics, or power cycles during the pre-check phase

#### Node1 or node2 crashes before the pre-check phase with HA still enabled

If either node1 or node2 crashes before the pre-check phase, no aggregates have been relocated yet and the HA configuration is still enabled.

#### About this task

Takeover and giveback can proceed normally.

#### Steps

1. Check the console for EMS messages that the system might have issued, and take the recommended corrective action.
2. Continue with the node-pair upgrade procedure.

### Reboots, panics, or power cycles during first resource-release phase

### Node1 crashes during the first resource-release phase with HA still enabled

Some or all aggregates have been relocated from node1 to node2, and HA is still enabled. Node2 takes over node1's root volume and any non-root aggregates that were not relocated.

#### About this task

Ownership of aggregates that were relocated look the same as the ownership of non-root aggregates that were taken over because the home owner has not changed.

When node1 enters the `waiting for giveback` state, node2 gives back all of the node1 non-root aggregates.

#### Steps

1. After node1 is booted up, all the non-root aggregates of node1 have moved back to node1. A manual aggregate relocation needs to be done of the aggregates from node1 to node2 using the command:

```
storage aggregate relocation start -node node1 -destination node2 -
aggregate-list * -ndocontroller-upgrade true
```

2. Continue with the node-pair upgrade procedure.

### Node1 crashes during the first resource-release phase while HA is disabled

Node2 does not take over but it is still serving data from all non-root aggregates.

#### Steps

1. Bring up node1.
2. Continue with the node-pair upgrade procedure.

### Node2 fails during the first resource-release phase with HA still enabled

Node1 has relocated some or all of its aggregates to node2. HA is enabled.

#### About this task

Node1 takes over all of node2's aggregates as well as any of its own aggregates that it had relocated to node2. When node2 boots up, the aggregate relocation is completed automatically.

#### Steps

1. Bring up node2.
2. Continue with the node-pair upgrade procedure.

### Node2 crashes during the first resource-release phase and after HA is disabled

Node1 does not take over.

#### Steps

1. Bring up node2.  
A client outage occurs for all aggregates while node2 is booting up.
2. Continue with rest of the node-pair upgrade procedure.

## Reboots, panics, or power cycles during the first verification phase

### Node2 crashes during the first verification phase with HA disabled

Node3 does not take over following a node2 crash as HA is already disabled.

#### Steps

1. Bring up node2.  
A client outage occurs for all aggregates while node2 is booting up.
2. Continue with the node-pair upgrade procedure.

### Node3 crashes during the first verification phase with HA disabled

Node2 does not take over but it is still serving data from all non-root aggregates.

#### Steps

1. Bring up node3.
2. Continue with the node-pair upgrade procedure.

## Reboots, panics, or power cycles during first resource-regain phase

### Node2 crashes during the first resource-regain phase during aggregate relocation

Node2 has relocated some or all its aggregates of node1 to node3. Node3 serves data from aggregates that were relocated. HA is disabled and hence there is no takeover.

#### About this task

There is client outage for aggregates that were not relocated. On booting up node2, the aggregates of node1 are relocated to node3.

#### Steps

1. Bring up node2.
2. Continue with the node-pair upgrade procedure.

### Node3 crashes during the first resource-regain phase during aggregate relocation

If node3 crashes while node2 is relocating aggregates to node3, the task continues after node3 boots up.

#### About this task

Node2 continues to serve remaining aggregates, but aggregates that were already relocated to node3 encounter client outage while node3 is booting up.

#### Steps

1. Bring up node3.
2. Continue with the controller upgrade.

## Reboots, panics, or power cycles during post-check phase

### Node2 or node3 crashes during the post-check phase

The HA pair is disabled hence this is no takeover. There is a client outage for aggregates belonging to the node that rebooted.

#### Steps

1. Bring up the node.
2. Continue with the node-pair upgrade procedure.

## Reboots, panics, or power cycles during second resource-release phase

### Node3 crashes during the second resource-release phase

If node3 crashes while node2 is relocating aggregates, the task continues after node3 boots up.

#### About this task

Node2 continues to serve remaining aggregates, but aggregates that were already relocated to node3 and node3's own aggregates encounter client outages while node3 is booting.

#### Steps

1. Bring up node3.
2. Continue with the controller upgrade procedure.

### Node2 crashes during the second resource-release phase

If node2 crashes during aggregate relocation, node2 is not taken over.

#### About this task

Node3 continues to serve the aggregates which it has been relocated but the aggregates owned by node2 encounter client outages.

#### Steps

1. Bring up node2.
2. Continue with the controller upgrade procedure.

## Reboots, panics, or power cycles during the second verification phase

### Node3 crashes during the second verification phase

If node3 crashes during this phase, takeover does not happen since HA is already disabled.

#### About this task

#### Steps

1. Bring up node3.

A client outage occurs for all aggregates while node3 is booting up.

2. Continue with the node-pair upgrade procedure.

### Node4 crashes during the second verification phase

If node4 crashes during this phase, takeover does not happen. Node3 serves data from the aggregates.

#### About this task

There is an outage for non-root aggregates that were already relocated until node4 reboots.

#### Steps

1. Bring up node4.
2. Continue with the node-pair upgrade procedure.

## Issues that can arise in multiple stages of the procedure

Some issues can occur during different stages of the procedure.

### Unexpected `storage failover show` command output

During the procedure, if the node that hosts all data aggregates panics or is rebooted accidentally, you might see unexpected output for the `storage failover show` command before and after the reboot, panic, or power cycle.

#### About this task

You might see unexpected output from the `storage failover show` command in Stage 2, Stage 3, Stage 4, or Stage 5.

The following example shows the expected output of the `storage failover show` command if there are no reboots or panics on the node that hosts all the data aggregates:

```
cluster::> storage failover show
Node           Partner           Takeover
-----
node1          node2             false      Unknown
node2          node1             false      Node owns partner aggregates
as part of the non-disruptive head upgrade procedure. Takeover is not
possible: Storage failover is disabled.

2 entries were displayed.
```

The following example shows the output of the `storage failover show` command after a reboot or panic:

```
cluster::> storage failover show
Node           Partner           Takeover
-----
node1          node2             -          Unknown
node2          node1             false     Waiting for node1, Partial
```

```
giveback, Takeover is not possible: Storage failover is disabled  
  
2 entries were displayed.
```

Although the output says that a node is in partial giveback and that storage failover is disabled, you can disregard this message.

**Step**

1. No action is required; continue with the node-pair upgrade procedure.

## LIF migration failure

After you migrate them, LIFs might not come online after migration in Stage 2, Stage 3, or Stage 5.

**Steps**

1. Ensure that port MTU size is the same as that of the source node.  
For example, if the cluster port MTU size is 9000 on the source node, it should be 9000 on the destination node.
2. Check the physical connectivity of the network cable if the physical state of the port is down.

## Copyright information

---

Copyright © 2018 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

## Trademark information

---

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

## How to send comments about documentation and receive update notifications

---

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[\*doccomments@netapp.com\*](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277



**M**

- manual steps [5](#)
- mapping ports from node1 to node3 [34](#)

**N**

- NAS data LIFs
  - moving from node2 to node3 [40](#)
- netboot
  - preparing nodes for [22](#)
- nodes
  - preparing for netboot [22](#)
  - verifying readiness before upgrading [13](#)
- non-root aggregates
  - relocating from node2 to node3 [43](#)
- nondisruptive controller update
  - controller upgrade simplification [5](#)
- nondisruptive hardware upgrades
  - how aggregate relocation (ARL) allows these [5](#)
- nondisruptive node-pair upgrades
  - guidelines for controllers using aggregate relocation [9](#)
  - how they work using aggregate relocation [10](#)

**O**

- onboard key managers
  - using to manage authentication [18](#)

**P**

- ports
  - mapping ports from node2 to node4 [56](#)
- postupgrade
  - decommissioning the old system [69](#)
- prechecks
  - correcting aggregate ownership if precheck fails [16](#)

**Q**

- quorum
  - ensuring node with different network ports has joined [37](#), [60](#)

**R**

- relationships
  - quiescing for SnapMirror [18](#)

**S**

- SAN LIFs
  - verifying presence on node3 [40](#)
  - when to delete [45](#)
- SnapMirror relationships

- quiescing [18](#)
- stage 4 of the process) [43](#)
- storage encryption
  - rekeying after upgrading [17](#)
- Storage Encryption
  - setting up on new controller module [68](#)
- suggestions
  - how to send feedback about documentation [80](#)

**T**

- troubleshooting
  - aggregate relocation failure [71](#)
  - checking for aggregate relocation failure [71](#)
  - failures during or after Stage 3) [74](#)
  - failures occur during Stage 5 [75](#)
  - identifying when in the process a failure occurs [72](#)
  - LIF migration failure [77](#)
  - node 3 crashes during aggregate relocation [74](#)
  - node crashes before Stage 2 [72](#)
  - node1 crashes during or after Stage 2 [73](#)
  - node2 crashes after Stage 3 [75](#)
  - node2 fails after relocating aggregates [74](#)
  - node2 fails after Stage 2 with HA disabled [73](#)
  - node2 fails during or after Stage 2 [73](#)
  - node2 fails to take over [73](#)
  - node2 fails to take over but still serves data [74](#)
  - node3 fails after Stage 3 [75](#)
  - node3 fails during Stage 5 [75](#)
  - node3 fails to take over during Stage 3) [74](#)
  - node4 fails during Stage 5 [76](#)
  - relocating aggregates after completing the upgrade [71](#)
  - unable to boot node3 after Stage 3 failure [75](#)
  - unexpected output for storage failover show command [76](#)
- Twitter
  - how to receive automatic notification of documentation changes [80](#)

**U**

- UTA/UTA2 card
  - setting if the node has a UTA/UTA2 card [27](#)
- UTA/UTA2 ports
  - checking and configuring [52](#)
  - checking and configuring ports that contain [29](#)
  - settings for [50](#)

**V**

- vetoed aggregates
  - correcting relocation of [21](#)
- VLANs
  - re-creating during aggregate relocation [39](#), [62](#)